NeuroImage 166 (2018) 385–399

Contents lists available at ScienceDirect

NeuroImage

journal homepage: www.elsevier.com/locate/neuroimage

Beyond modularity: Fine-scale mechanisms and rules for brain network reconfiguration

Ankit N. Khambhati [a], Marcelo G. Mattar [a,b], Nicholas F. Wymbs [c], Scott T. Grafton [d], Danielle S. Bassett [a,e,*]

[a] Department of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA
[b] Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104, USA
[c] Department of Physical Medicine and Rehabilitation, Johns Hopkins Medical Institution, Baltimore, MD 21205, USA
[d] Department of Psychological and Brain Sciences, University of California, Santa Barbara, CA 93106, USA
[e] Department of Electrical and Systems Engineering, University of Pennsylvania, Philadelphia, PA 19104 USA

ARTICLE INFO

Keywords:
Network neuroscience
Non-negative matrix factorization
Community detection
Subgraph
Cognitive control
Functional connectivity

ABSTRACT

The human brain is in constant flux, as distinct areas engage in transient communication to support basic behaviors as well as complex cognition. The collection of interactions between cortical and subcortical areas forms a functional brain network whose topology evolves with time. Despite the nontrivial dynamics that are germane to this networked system, experimental evidence demonstrates that functional interactions organize into putative brain systems that facilitate different facets of cognitive computation. We hypothesize that such dynamic functional networks are organized around a set of rules that constrain their spatial architecture – which brain regions may functionally interact – and their temporal architecture – how these interactions fluctuate over time. To objectively uncover these organizing principles, we apply an unsupervised machine learning approach called non-negative matrix factorization to time-evolving, resting state functional networks in 20 healthy subjects. This machine learning approach automatically groups temporally co-varying functional interactions into subgraphs that represent putative topological modes of dynamic functional architecture. We find that subgraphs are stratified based on both the underlying modular organization and the topographical distance of their strongest interactions: while many subgraphs are largely contained within modules, others span between modules and are expressed differently over time. The relationship between dynamic subgraphs and modular architecture is further highlighted by the ability of time-varying subgraph expression to explain inter-individual differences in module reorganization. Collectively, these results point to the critical role that subgraphs play in constraining the topography and topology of functional brain networks. More broadly, this machine learning approach opens a new door for understanding the architecture of dynamic functional networks during both task and rest states, and for probing alterations of that architecture in disease.

Introduction

More than just a sum of its parts, the brain performs computations and processes information by linking functionally specialized areas through complex patterns of anatomical wiring (Alexander et al., 1990; Felleman and Van Essen, 1991; Van Essen et al., 1992; Sporns et al., 2005). Indeed, the underlying structural network forms the foundation of a wide repertoire of functional interactions between different regions (Deco et al., 2012; Senden et al., 2012; Deco and Jirsa, 2012). Collectively, these interactions can be modeled as edges between nodes in a graph (Bassett and Bullmore, 2006; Reijneveld et al., 2007; Bullmore and Sporns, 2009; Bullmore and Bassett, 2011; Kaiser, 2011; Sporns, 2013) to probe the neurophysiological underpinnings of thought, perception, and action (Sporns, 2014; Medaglia et al., 2015). Importantly, to actuate behavior and cognition through a changing landscape of environmental demands, these patterns of functional interactions must flexibly reconfigure (Hutchison et al., 2013; Calhoun et al., 2014; Kopell et al., 2014; Mattar et al., 2015), presumably according to organizing principles that coordinate the dynamic engagement and disengagement of distinct sets of brain areas (Bassett et al., 2011; Doron et al., 2012; Park and Friston, 2013; Chai et al., 2016b).

A fundamental core of this dynamic architecture is thought to be

* Corresponding author. Department of Bioengineering, University of Pennsylvania, Philadelphia, PA 19104, USA.
E-mail address: dsb@seas.upenn.edu (D.S. Bassett).

https://doi.org/10.1016/j.neuroimage.2017.11.015
Received 5 September 2017; Received in revised form 18 October 2017; Accepted 7 November 2017
Available online 11 November 2017

modularity – the division of functionally engaged brain regions into putative modules that may compartmentalize computation within discrete functional systems – such as motor, visual, auditory, or attention – without disturbing brain regions in other systems (Newman, 2006; Sporns and Betzel, 2016). Specifically, modules represent discrete clusters of a graph in which nodes of the same module are more strongly interconnected to one another than to nodes of different modules. The composition of a module may change with time as the edges in the network reconfigure over time. Reconfiguration of modules in brain networks is thought to support functional dynamics driving behavior and cognition by compartmentalizing integrated and segregated neural processing of individual brain regions (Bertolero et al., 2015; Mattar et al., 2015). Moreover, functional brain networks exhibit flexibility in their module composition as they adapt to cognitive demands associated with completing a task (Braun et al., 2015, 2016; Mattar et al., 2015; Telesford et al., 2016), processing linguistic stimuli (Doron et al., 2012; Chai et al., 2016b), or learning a new skill (Bassett et al., 2011, 2013b, 2015). Notably, individual differences in flexibility are correlated with individual differences in learning (Bassett et al., 2011; Gerraty et al., 2016), working memory performance (Braun et al., 2015), attention (Shine et al., 2016), and cognitive flexibility (Braun et al., 2015), which is particularly interesting in light of its alteration in aging (Schlesinger et al., 2017) and depression (Wei et al., 2017), and its role as an intermediate phenotype in schizophrenia (Braun et al., 2016).

Yet, while flexibility appears to be an important attribute of functional brain networks, a fundamental understanding of how network reconfiguration occurs, and what rules constrain the types of reconfiguration that characterize neural systems, is lacking. Do brain regions spontaneously arrange themselves into efficient modular configurations, or do functional interactions obey a distinct set of rules that constrain which modules can and cannot exist? Are there separable groups of functional interactions that differentially drive integration and segregation among network modules? While modules partition groups of brain regions based purely on the presence or strength of functional interactions, they remain agnostic to the *topology* of the functional interactions within and between them. These open questions have motivated the development of a new wave of graph theoretic tools grounded in machine learning that recover latent structure in dynamic brain networks as coherent groups of temporally co-varying functional interactions – known as subgraphs.

Mathematically, functional subgraphs form a basis set of unique patterns of graph edges whose weighted linear combination – given by a set of time-varying basis weights for each subgraph – reconstructs a repertoire of graph configurations observed over time. All graph nodes participate to a varying degree in each subgraph, and the edges between nodes are assigned weights based on how strongly they co-vary over time. Thus, subgraphs are computationally represented as weighted adjacency matrices of the same size as the original graph – enabling us to query subgraph architecture as we might with brain graphs. These properties underlie the important distinction between modules and subgraphs insofar as modules represent hard-partitions of a graph's nodes into distinct clusters and subgraphs represent soft-partitions of a graph's edges into a set of overlapping and non-orthogonal basis components – enabling a node to participate in multiple contexts within the graph. More conceptually, functional subgraphs can be thought of as topological modes of interacting brain regions that are differentially expressed over time (Leonardi et al., 2013; Chai et al., 2016a). Spatially, individual brain regions may engage and play distinct topological roles among different groups of brain regions in multiple subgraphs – a capability that is critical for capturing network architecture of putative functional brain systems (Chai et al., 2016a).

The framework of functional subgraphs yields an opportunity to examine subgraph topology and subgraph dynamics associated with different cognitive states and different behaviors. For example, functional brain networks could decompose into different components of visual processing in which visual areas interact with ventral attention areas in one subgraph and with dorsal attention areas in another – with each subgraph increasing or decreasing its relative expression during different phases of cognitive processing. A recent application of the subgraph decomposition technique in neurodevelopment demonstrated that functional brain networks of children and young adults are composed of subgraphs representing known brain systems that are common across both age groups yet differ in the temporal properties of their expression during the resting state (Chai et al., 2016a). Such a decomposition has also been used to uncover putative network subregions in epilepsy that play unique roles in the initiation and maintenance of neural dysfunction (Khambhati et al., 2017). Despite the apparent role of subgraphs as functional substrates of information processing in the brain, their topographical and topological properties are not well understood. Do subgraphs reveal network architecture that is spatially distributed across the brain? Are subgraphs bound to the modular architecture of brain networks or do they span between modules?

Here, we develop a quantitative framework for identifying functional subgraphs and characterizing their relationship to whole brain network architecture. We base our analysis on a burgeoning application of non-negative matrix factorization (NMF) that decomposes dynamic functional networks into constituent additive parts rather than generalized features (Lee et al., 1999). In implementing NMF in the context of neuroimaging data, we detail several important methodological considerations. Our flexible approach enables subgraph analysis across multiple spatiotemporal scales of network topology and dynamics through manipulation of three main parameters. By optimizing these parameters, we uncover a robust set of subgraphs across multiple human subjects and examine their sensitivity to functional interactions within and between network modules, and over different geographic distances.

The topological and topographical sensitivity of subgraphs – a basis set of temporally co-varying network edges – would highlight important constraints on the temporal organization of dynamical brain networks. Based on prior work suggesting distance-dependent organization of brain networks into local, function-specific interactions (characteristic of clusters and modules (Rubinov and Sporns, 2010; Telesford et al., 2011)) and distributed, integrative interactions (characteristic of hubs and rich-clubs (Bertolero et al., 2017; Senden et al., 2014; de Reus and van den Heuvel, 2013; van den Heuvel and Sporns, 2011)), we first hypothesize that subgraphs are selectively sensitive to functional interactions over different distances. Specifically, we hypothesize that short-range interactions are more likely to exist in some subgraphs and long-range interactions are more likely to exist in other subgraphs. Critically, distance-based rules for subgraph organization would support a theory that functional interactions extending over similar geographical distances are more likely to co-vary during resting-state processes. Second, we expect that subgraphs are differentially sensitive to functional interactions within modules and functional interactions between modules. Functional interactions within the same module may show more similar patterns of temporal variation – and are therefore more likely to exhibit greater weight in one set of subgraphs – than functional interactions that span between modules – which are therefore more likely to exhibit greater weight in another set of subgraphs. Subgraphs sensitive to network topology within modules might also exhibit a strong correlation between fluctuation in their temporal expression and flexibility of module reorganization over time. Such a relationship would provide a novel perspective on the inter-regional changes of functional interactions that accompany meso-scale alterations in functional brain networks in both health and disease (Braun et al., 2016; Rubinov and Bassett, 2011; Fornito et al., 2015).

Methods

Experimental design

Main dataset

Twenty participants (nine female; ages 19–53 years; mean age = 26.7

years) with normal or corrected vision and no history of neurological disease or psychiatric disorders were recruited for this experiment, non-overlapping results from which have been reported elsewhere (Ashourvan et al., 2017; Mattar et al., 2016a, 2016b; Gu et al., 2017). Each participant underwent four separate, ten minute, resting state scanning sessions on consecutive days. All participants volunteered and provided informed consent in writing in accordance with the guidelines of the Institutional Review Board of the University of Pennsylvania (IRB #801929).

Validation dataset

In accordance with the guidelines set out by the Institutional Review Board of the University of California, Santa Barbara, twenty-two right-handed participants (13 females and 9 males) volunteered to participate and provided informed consent in writing. Two participants were excluded from the following analyses: one failed to complete the entirety of the experiment and the other had persistent head motion greater than 5 mm during MRI scanning. Each participant underwent four separate, eight minute, resting state scanning sessions on separate days. Separate analyses of the data acquired in this study are reported elsewhere (Bassett et al., 2013a, 2015; Wymbs and Grafton, 2014).

Data acquisition and pre-processing

Main dataset

Magnetic resonance images were obtained at the Hospital of the University of Pennsylvania using a 3.0 T Siemens Trio MRI scanner equipped with a 32-channel head coil. T1-weighted structural images of the whole brain were acquired on the first of four scan sessions per subject using a three-dimensional magnetization-prepared rapid acquisition gradient echo pulse sequence (repetition time (TR) 1620 ms; echo time (TE) 3.09 ms; inversion time 950 ms; voxel size 1 mm × 1 mm × 1 mm; matrix size 190 × 263 × 165). A field map was also acquired at each of the four scan sessions (TR 1200 ms; TE1 4.06 ms; TE2 6.52 ms; flip angle 60 °; voxel size 3.4 mm × 3.4 mm × 4.0 mm; field of view 220 mm; matrix size 64 × 64 × 52) to correct geometric distortion caused by magnetic field inhomogeneity. In all scans, T2*-weighted images sensitive to blood oxygenation level-dependent contrasts were acquired using a slice accelerated multiband echo planar pulse sequence (TR 500 ms; TE 30 ms; flip angle 30°; voxel size 3.0 mm × 3.0 mm × 3.0 mm; field of view 192 mm; matrix size 64 × 64 × 48).

Validation dataset

Magnetic resonance images were acquired using a 3.0 T Siemens Trio with a 12-channel phased-array head coil. Each whole-brain scan epoch was created using a single-shot echo planar imaging sequence that was sensitive to BOLD contrast to acquire 37 slices per repetition time (repetition time (TR) of 2000 ms, 3 mm thickness, 0.5 mm gap) with an echo time of 30 ms, a flip angle of 90°, a field of view of 192 mm, and a 64 × 64 acquisition matrix. Before the first round of data collection, we acquired a high-resolution T1-weighted sagittal sequence image of the whole brain (TR of 15.0 ms, echo time of 4.2 ms, flip angle of 90°, 3D acquisition, field of view of 256 mm, slice thickness of 0.89 mm, and 256 × 256 acquisition matrix).

fMRI preprocessing

We preprocessed the resting state fMRI data using FEAT (FMRI Expert Analysis Tool) Version 6.00, part of FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl). Specifically, we applied: EPI distortion correction using FUGUE (Jenkinson, 2004); motion correction using MCFLIRT (Jenkinson et al., 2002); slice-timing correction using Fourier-space timeseries phase-shifting; non-brain removal using BET (Smith, 2002); grand-mean intensity normalization of the entire 4D dataset by a single multiplicative factor; and highpass temporal filtering (Gaussian-weighted least-squares straight line fitting, with sigma = 50.0s).

Nuisance time-series were voxelwise regressed from the preprocessed data. Nuisance regressors included (i) three translation (X, Y, Z) and three rotation (Pitch, Yaw, Roll) timeseries derived by retrospective head motion correction ($R = [X, Y, Z, pitch, yaw, roll]$), together with expansion terms ($[RR^2R_{t-1}R^2_{t-1}]$), for a total of 24 motion regressors (Friston et al., 1996)); (ii) the first five principal components calculated from timeseries derived from regions of non-interest (white matter and cerebrospinal fluid), using the anatomical CompCor method (aCompCor) (Behzadi et al., 2007) and (iii) the average signal derived from white matter voxels located within a 15 mm radius from each voxel, following the ANATICOR method (Jo et al., 2010). Global signal was not regressed out of voxel time series (Murphy et al., 2009; Saad et al., 2012; Chai et al., 2012). Finally, the mean functional image and the Harvard-Oxford atlas were co-registered using Statistical Parametric Mapping software (SPM12; Wellcome Department of Imaging Neuroscience, www.fil.ion.ucl.ac.uk/spm) in order to extract regional mean time-series (Fig. 1A). We separately co-registered the mean functional image and a multi-modal cortical parcellation (Glasser et al., 2016) to test the robustness of our analysis to the choice of parcellation (see Supplemental Information).

Constructing time-varying functional networks

To construct dynamic functional brain networks, we begin by dividing the BOLD signal over the four scan sessions into 80 non-overlapping time windows – each 30 s in duration and containing spectral information between the frequencies of 0.03–1.00 Hz (Fig. 1B). For each time window, we measure functional interactions between each pair of brain regions based on coherence within a frequency band of 0.03–0.20 Hz. We use the *mtspec* Python implementation of multi-taper coherence (Prieto et al., 2009) with time-bandwidth product of 2.5 and 4 tapers to achieve a frequency resolution of 0.08 Hz. Coherence values are stored in a time-varying, $N \times N \times T$ adjacency matrix $\mathbf{A}$, where $N = 112$ brain regions and $T = 80$ time windows, for each subject (Fig. 1C). In our weighted network analysis, we retain and analyze *all* functional interactions between brain regions, and do not apply any threshold or perform any binarization.

An alternate representation of the three-dimensional network adjacency matrix $\mathbf{A}$ is a two-dimensional network configuration matrix $\widehat{\mathbf{A}}$, which tabulates all $N \times N$ pairwise interactions across T time windows (Fig. 1D, **right**). Due to symmetry of the adjacency matrix for every time window $t \in T$, we unravel the upper triangle of $\mathbf{A}_t$, resulting in the weights of $E = \frac{N(N-1)}{2}$ functional interactions. Thus, $\widehat{\mathbf{A}}$ has dimensions $E \times T$. We construct a separate network configuration matrix for each subject.

To test that the resulting resting state functional networks are indeed representative of dynamic and time-varying functional interactions, in line with previous studies (Zalesky et al., 2014a; Hindriks et al., 2016), we compare the variance of dynamic functional interactions to the variance of static functional interactions of surrogate data (see Supplemental Information for methodological details). By conducting a paired *t*-test between the distribution of variances for static, surrogate functional interactions to the distribution of variances for dynamic, true functional interactions, we find that indeed the variance of dynamic functional interactions is greater than the variance of static functional interactions (Fig. S1; $t_{19} = 13.6$, $p = 3.0 \times 10^{-11}$). Thus, our dynamic functional brain networks exhibit greater fluctuations in their interactions than expected by a static functional network – providing evidence that we are indeed capturing dynamic, time-varying functional interactions.

Partitioning the dynamic functional network

Clustering nodes into modules

To identify functional modules – partitions of highly inter-connected brain regions – we maximize the following multi-layer modularity quality function $\mathscr{Q}$ that assigns brain regions into modules that vary with

time (Mucha et al., 2010):

$$\mathscr{Q} = \frac{1}{2\mu}\sum_{ijt_0t_1}\left[\left(A_{ijt_0} - \gamma_{t_0}P_{ijt_0}\right)\delta_{t_0t_1} + \delta_{ij}\omega_{jt_0t_1}\right]\delta\left(g_{it_0}, g_{jt_1}\right), \quad (1)$$

where indices i, j represent brain regions and t_0, t_1 represent consecutive time windows, μ is the sum of all functional interactions in the dynamic network, $\mathbf{P}$ represents functional interactions derived from a null model (e.g., the Newman-Girvan null model (Newman and Girvan, 2004)), γ is a structural resolution parameter for modular organization within a single time window, ω is a temporal resolution parameter to model changes in modular organization over time, and g is the module assignment of a brain region within a time window (Bassett et al., 2013a). In accord with prior studies, we choose resolution parameters such that $\gamma_{t_0} = \gamma = 1$ and $\omega_{jt_0t_1} = \omega = 1$ (Bassett et al., 2011, 2013b; Braun et al., 2015). To identify functional modules common to all subjects, we linked time windows between subjects using an inter-subject resolution parameter $\rho = 1.0/T$, where T is the number of time windows.

We use a Louvain-like locally greedy optimization to maximize $\mathscr{Q}$ (Jeub et al., 2016). The optimization landscape of the multi-layer modularity contains a plateau with many near-optimal solutions (Good et al., 2010). This near-degeneracy can be addressed by aggregating solutions over several runs of the Louvain algorithm (Bassett et al., 2013b). We separately perform 100 independent optimizations of $\mathscr{Q}$ and aggregate the probability of each pair of brain regions being co-assigned to the same functional module across all time windows, subjects, and optimization runs into a single module allegiance matrix of size $N \times N$ (Bassett et al., 2015).

To form a consensus partition of brain regions into functional modules, we applied a single-layer formulation of the Louvain-like locally greedy optimization of the modularity quality function to the module allegiance matrix using an optimum $\gamma \in [0.5, 3.0]$ that yields an elbow in $\mathscr{Q}$ (Khambhati et al., 2015).

Decomposing dynamic network edges into subgraphs

To identify functional subgraphs – a basis set of temporally co-varying network interactions – we apply an unsupervised machine learning algorithm called non-negative matrix factorization (NMF) (Lee et al., 1999) to the network configuration matrix (Fig. 1D). NMF decomposes the network configuration matrix into constituent subgraphs and accompanying time-varying expression coefficients (Chai et al., 2016a; Khambhati et al., 2017). Each subgraph is an additive component of the original network – weighted by its associated time-varying expression coefficient – and represents a pattern of functional interactions between brain regions. The NMF-based subgraph learning paradigm is a basis decomposition of a collection of dynamic graphs that separates co-varying network edges into subgraphs – or basis functions – with associated temporal coefficients – or basis weights. Unlike other graph clustering approaches that seek a hard partition of nodes and edges into clusters (Mucha et al., 2010; Bassett et al., 2013a), the temporal coefficients provide a soft partition of the network edges, such that the original functional network of any time window can be reconstructed through a linear combination of all of the subgraphs weighted by their associated temporal coefficient in that time window (Leonardi et al., 2013, 2014; Chai et al., 2016a; Khambhati et al., 2017). This computation implies that at a specific time window, subgraphs with a high temporal coefficient contribute their pattern of functional interactions more than subgraphs with a low temporal coefficient.

Mathematically, NMF approximates $\widehat{\mathbf{A}}$ by two non-negative matrices $\mathbf{W}$ – the subgraph matrix identifying patterns of functional interactions (with dimensions $E \times m$) – and $\mathbf{H}$ – the time-varying expression coefficients matrix (with dimensions $m \times T$) – such that:

$$\min_{\mathbf{W},\mathbf{H}} \ \frac{1}{2}\| \widehat{\mathbf{A}} - \mathbf{WH}\|_F^2 + \alpha \| \mathbf{W}\|_F^2 + \beta\sum_{t=1}^{T} \| H(:,t)\|_1^2, \quad (2)$$

where $m \in [2,\ \min(E,T) - 1]$ is the number of subgraphs to decompose, β is a penalty weight to impose sparse temporal expression coefficients, and α is a regularization of the interaction strengths for subgraphs (Kim and Park, 2011). To solve the NMF equation, we use an alternating non-negative least squares with block-pivoting method with 100 iterations for fast and efficient factorization of large matrices (Kim et al., 2014). We initialized $\mathbf{W}$ and $\mathbf{H}$ with non-negative weights drawn from a uniform random distribution on the interval $[0, 1]$.

To select the parameters m, β, and α, we pursue a random sampling scheme – shown to be effective in optimizing high-dimensional parameter spaces (Bergstra and Bengio, 2012) – in which we re-run the NMF algorithm for 5000 parameter sets in which m is drawn from $\mathscr{U}(2, 30)$, β is drawn from $\mathscr{U}(0.01, 1)$, and α is drawn from $\mathscr{U}(0.01, 1)$ (Fig. 2). We evaluate subgraph learning performance based on a 10-fold cross-validation scheme for each parameter set and compute the following three quality measures: cross-validation error on held-out fold ($\widehat{\mathbf{A}} - \mathbf{WH}_F^2$), temporal sparsity of the held-out fold ($\frac{1}{T^*m}\sum_{t\in T}\sum_{i\in m}[H_{t,i} = 0]$), and subgraph sparsity of the training folds ($\frac{1}{E^*m}\sum_{e\in E}\sum_{i\in m}[W_{e,i} = 0]$). The optimal parameter set should yield subgraphs that minimize the cross-validation error and reliably span the space of observed network topologies of the held-out fold. Based on these criteria, we identified an optimum parameter set $(\overline{m}, \overline{\beta}, \overline{\alpha})$ that exhibits a low residual error in the bottom 25th percentile of our random sampling scheme.

Due to the non-deterministic nature of this approach, we integrated subgraph estimates over multiple runs of the algorithm using *consensus clustering* – a general method of testing robustness and stability of clusters over many runs of one or more non-deterministic clustering algorithms (Monti et al., 2003). Our adapted consensus clustering procedure (Greene et al., 2008; Greene, 2009) entailed the following steps: (i) run the NMF algorithm R times per network configuration matrix, (ii) concatenate subgraph matrix $\mathbf{W}$ across R runs into an aggregate matrix with dimensions $E \times (R^*\overline{m})$, and (iii) apply NMF to the aggregate matrix to determine a final set of subgraphs and expression coefficients.

Test-retest reliability of subgraphs. It is important to consider the reproducibility of subgraphs measured using NMF given new data. To quantify the reproducibility of functional subgraphs, we measured the extent to which the pattern of subgraph edge weights measured in one dataset predicts the pattern of subgraph edge weights measured in a second dataset. Specifically, we first applied NMF using the optimal parameter set to two different datasets ($\widehat{\mathbf{A}}_1$ and $\widehat{\mathbf{A}}_2$), resulting in two subgraph matrices ($\mathbf{W}_1$ and $\mathbf{W}_2$). Note that the subgraphs along the columns of $\mathbf{W}_1$ may not necessarily be ordered similarly to the subgraphs along the columns of $\mathbf{W}_2$ due to the stochastic nature of the NMF algorithm. To reorder subgraphs from the second dataset such that they correspond to the same order as subgraphs from the first dataset, we sought a mapping $X_{i,j}$ of subgraph W_1^i to subgraph W_2^j, where X is a Boolean matrix that prescribes whether the i^{th} subgraph from the first dataset is uniquely assigned to the j^{th} subgraph from the second dataset. The cost $C_{i,j}$ associated with assigning W_1^i to W_2^j is equal to $\left\|W_1^i - W_2^j\right\|$. To determine a unique X, we minimized the cost function $\sum_i\sum_j C_{i,j}X_{i,j}$ using the well-known Hungarian algorithm (Kuhn, 1955). After calculating an optimal assignment between subgraphs of the two datasets, we measured the similarity in the pattern of edge weights between assigned subgraph pairs (i, j) by computing the Pearson correlation coefficient. This approach enabled us to assess the reproducibility of each individual subgraph based on the magnitude of the Pearson correlation similarity measure relative to that expected by chance.

Geometrically-unconstrained surrogate subgraph model. Network geometry plays an important role in the organization of brain networks – with nearby pairs of brain regions more likely to exhibit strong interactions

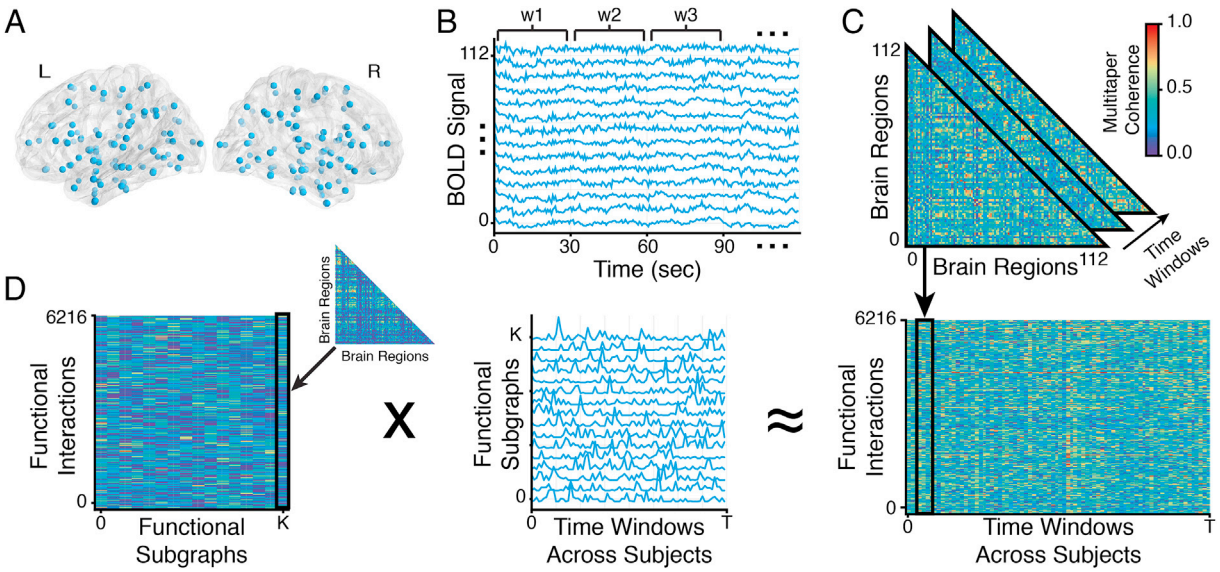

Fig. 1. Subgraphs of dynamic functional brain networks. (*A*) We obtained multi-band fMRI BOLD signals from 112 functional regions of interest (HO-112 atlas) in cortical and subcortical areas during the resting state from 20 healthy subjects across four separate imaging sessions. We constructed dynamic functional networks for each subject by (*B*) dividing the regional BOLD signal into 80 non-overlapping time windows, each 30 s (60 TRs) in duration, and (*C*) computing multitaper coherence between each pair of regional BOLD signals for every time window to obtain a time-varying adjacency matrix – with brain regions represented as network nodes and time-varying inter-regional coherence represented as weighted network edges. (*D*) We concatenated all pairwise edges over time and subjects to generate a single time-varying network configuration matrix for the entire sample of subjects (*right*). We applied non-negative matrix factorization (NMF) – which pursues a parts-based decomposition of the dynamic network – to the configuration matrix and clustered temporally co-varying network edges into a matrix of subgraphs (*left*) and a matrix of time-dependent coefficients (*middle*), which quantifies the level of expression in each time window for each subgraph.

than more distant pairs of brain regions (Roberts et al., 2016; Betzel et al., 2016). While the relationship between geometry and connectivity has been explored in static functional networks (Santarnecchi et al., 2014; Alexander-Bloch et al., 2012), their relationship in dynamic functional networks remains elusive. For instance, do functional interactions between nearby brain regions co-evolve differently than interactions between more distant brain regions? If so, NMF is more likely to separate interactions between nearby brain regions from interactions between distant brain regions into distinct subgraphs. In this study, we assessed whether a subgraph expresses functional interactions that follow geometric rules using a *geometrically-unconstrained* surrogate subgraph model. This surrogate model enabled us to measure the extent to which a subgraph exhibits geometric rules by chance, after the effect of geometry is removed from the dynamic functional network. Specifically, we first removed the effect of physical distance from the functional interaction weight in each time window by fitting a third-order polynomial between functional interaction strength and Euclidean distance of the adjoining brain regions and computing the residuals of the fit (details in (Roberts et al., 2016)). We next ran NMF on the residual dynamic functional interactions 1000 times using randomly generated initializations drawn from a uniform distribution. The resulting geometrically-unconstrained surrogate subgraphs represent groups of functional interactions that co-vary in strength over time with the effect of physical distance removed. We used these surrogate subgraphs to assess the extent to which subgraphs extract geometric relationships in dynamic networks.

Topologically-unconstrained surrogate subgraph model. In this study, we tested the hypothesis that subgraphs capture functional interactions that co-vary differently over time based on whether adjoining brain regions are located in the same functional module. When measuring the topological characteristics of different subgraphs, it is helpful to establish a null distribution of the measurement when meaningful topological architecture is removed from the subgraphs. In developing an appropriate surrogate model, we considered rewiring the network edges while accounting for the statistical properties of the edge weight distribution (Bassett et al., 2013a) and the geometrical properties of the edge weight variability over different physical distances (Roberts et al., 2016). Thus, we used a *topologically-unconstrained* surrogate subgraph model in which we randomly rewired the functional interactions between brain regions while preserving the geometric rules that govern the relationship between functional interaction strength and physical distance. Specifically, we fit a third-order polynomial between functional interaction strength and Euclidean distance, computed residual interaction strength after the fit, randomly permuted the residual interaction strengths between pairs of brain regions, and added the fitted geometric model back to the permuted, residual interaction strengths (details in (Roberts et al., 2016)). We next ran NMF on the resulting dynamic functional network 1000 times using randomly generated initializations drawn from a uniform distribution. The resulting topologically-unconstrained surrogate subgraphs represent groups of functional interactions that co-vary in strength over time with edge weights randomly distributed across the network according to geometric rules. We used these surrogate subgraphs to assess topological characteristics of subgraphs in which topological information is known to have been removed.

Measures of subgraph topology and dynamics

To assess whether the functional interactions of a subgraph are constrained by network topography, we compute the Pearson correlation coefficient between the strength of functional interactions and the Euclidean distance between their associated brain regions, for each subgraph (Fig. 3). Positive correlations imply stronger functional interactions over longer distances and negative correlations imply stronger functional interactions over shorter distances. We separately analyze correlations among inter-hemispheric and intra-hemispheric interactions to account for naturally longer inter-hemispheric distances compared to shorter intra-hemispheric distances. To test whether a subgraph is significantly sensitive to network topography, we compare the distance-wise Pearson correlations between the true subgraphs and the surrogate subgraphs derived from the geometrically-unconstrained model.

To determine whether the functional interactions of a subgraph are expressed within modules or between modules, we compute a module sensitivity index that maps the module allegiance of network regions onto subgraphs (Fig. 4). Mathematically, the module sensitivity index is

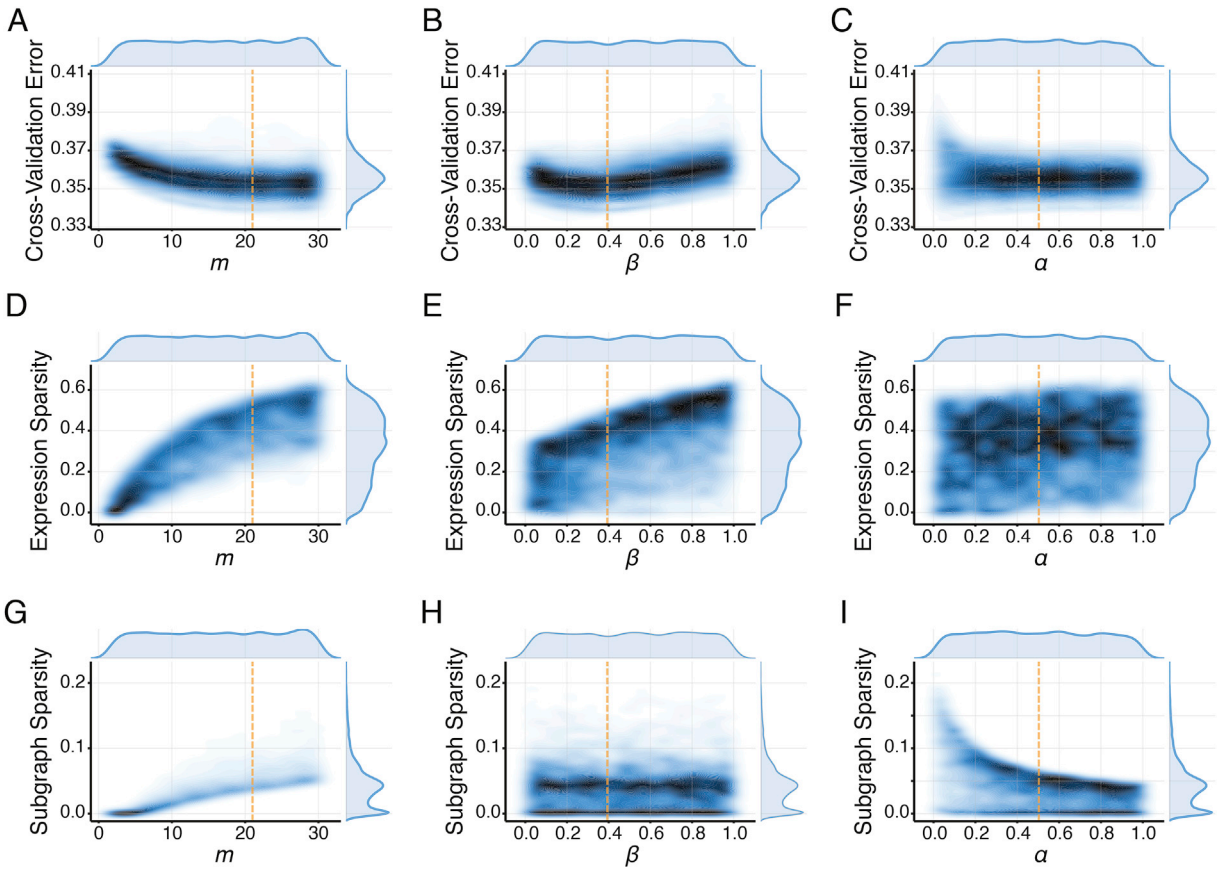

Fig. 2. Parameter optimization for subgraph detection. NMF-based subgraph detection requires optimizing three parameters: the number of subgraphs m, the temporal sparsity of subgraph expression β, and the regularization of subgraph edge weights α. To characterize this parameter space, we randomly sampled m, β, and α from a three-dimensional uniform distribution ($m \in [2,\ 30]$, $\beta \in [0.01,\ 1.0]$, $\alpha \in [0.01,\ 1.0]$) and applied NMF to each subject using each parameter set. Measurements associated with the optimal parameters are indicated by the dashed orange line. (*A-C*) The joint probability distribution of cross-validation error and marginalized m (Pearson $r = -0.42$, $p < 1 \times 10^{-16}$), marginalized β (Pearson $r = 0.36$, $p < 1 \times 10^{-16}$), and marginalized α (Pearson $r = 0.16$, $p < 1 \times 10^{-16}$) suggests that (i) increasing the number of subgraphs significantly improves the reconstruction of the original dynamic network, and (ii) increasing the subgraph expression sparsity or edge weight regularization slightly improves the reconstruction. (*D-F*) The joint probability distribution of percent sparse expression coefficients and marginalized m (Pearson $r = 0.79$, $p < 1 \times 10^{-16}$), marginalized β (Pearson $r = 0.52$, $p < 1 \times 10^{-16}$), and marginalized α (Pearson $r = 0.15$, $p < 1 \times 10^{-16}$) suggests that (i) increasing the number of subgraphs significantly introduces sparse temporal expression coefficients, and (ii) increasing the subgraph expression sparsity significantly increases the sparsity of the temporal expression coefficients, and that increasing the edge weight regularization has minimal effect on temporal expression sparsity. (*G-I*) The joint probability distribution of percent sparse subgraph edge weights and marginalized m (Pearson $r = 0.16$, $p < 1 \times 10^{-16}$), marginalized β (Pearson $r = -0.01$, $p < 1 \times 10^{-16}$), and marginalized α (Pearson $r = -0.53$, $p < 1 \times 10^{-16}$) suggests that increasing the number of subgraphs moderately increases the sparsity of subgraph edge weights and increasing edge weight regularization parameter strongly reduces the sparsity of subgraph edge weights.

computed for each subgraph as the Pearson correlation between the upper triangle of the module allegiance matrix and the upper triangle of the subgraph adjacency matrix. Positive correlations imply that stronger functional interactions are more likely to be found between brain regions that have a high probability of being organized in the same functional module; in contrast, negative correlations imply that stronger functional interactions are more likely to be found between brain regions that have a low probability of being organized in the same functional module. To test whether a subgraph is significantly sensitive to modular architecture, we compare the module sensitivity index between true subgraphs and surrogate subgraphs derived from the topologically-unconstrained model.

To link subgraph dynamics to the modular organization of brain networks, we compute an inter-subgraph coupling metric defined as the Pearson correlation between the expression coefficients of each possible pair of subgraphs (Fig. 5). Intuitively, pairs of subgraphs with more positive correlation exhibit greater coupling in their relative expression and pairs of subgraphs with more negative correlation exhibit lower coupling in their relative expression. We compare inter-individual differences in inter-subgraph coupling to network flexibility – a measure of how frequently an individual's network reorganizes its modular structure. Specifically, network flexibility is defined as the mean fraction of times that brain regions change module assignment in each session (Bassett et al., 2011).

To measure the extent to which a brain region participates in subgraph coupling, we computed the sum of a brain region's node strength in each subgraph, weighted by the average coupling of that subgraph. Brain regions with larger coupling coefficients are more influential participants in subgraphs that are, on average, more strongly coupled to all other subgraphs. We related this regional coupling to the regional flexibility measure – the mean fraction of times that a single brain region changes module assignment.

Results

Selecting parameters for NMF-based subgraph decomposition

Non-negative matrix factorization (NMF) requires parameter optimization that is critical for identifying a robust set of subgraphs that (i) best reconstruct the original dynamic network, and (ii) reflect functional interactions that are gradually expressed and not overly specific to a single point in time. We apply a random sampling scheme to characterize the rich parameter space of the number of subgraphs m, the temporal sparsity of subgraph expression β, and the regularization of subgraph

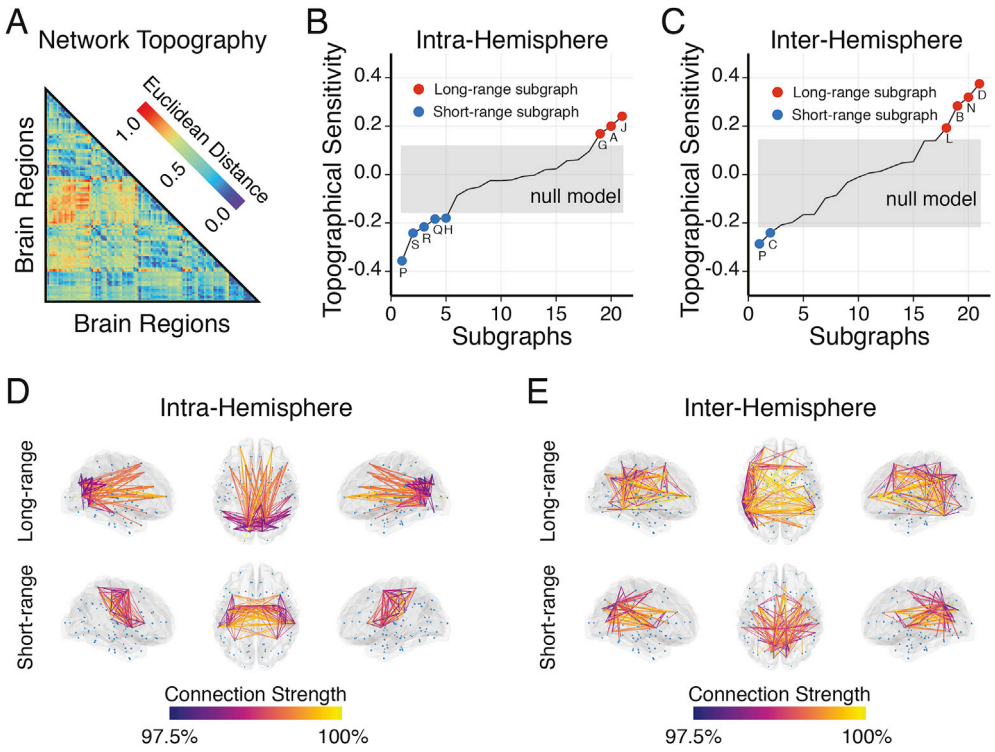

Fig. 3. Topographical properties of subgraphs. (*A*) Topographic map derived from the Harvard-Oxford anatomical atlas composed of 112 areas in cortex, subcortex, and brainstem, representing the Euclidean distance between pairs of brain regions in physical space. We quantified the sensitivity of each subgraph to topographic relationships in the map by computing the topographical sensitivity – defined as the Pearson correlation coefficient between the topographic distance and subgraph-specific edge strength between all pairs of brain regions (see Methods). A more positive correlation indicates that edge strength increases with increasing distance between brain regions, whereas a more negative correlation indicates that edge strength increases with decreasing distance between brain regions. To identify subgraphs with significant topographical sensitivity, we generated a set of geometrically-unconstrained surrogate subgraphs by applying NMF to dynamic functional networks in which distance-wise relationships were regressed out. (*B*) Distribution of intra-hemispheric topographical sensitivity for each subgraph – rank ordered from smallest to largest. Compared to the null model, subgraphs 1–5 (blue) exhibited significantly negative topographical sensitivity and subgraphs 19–21 (red) exhibited significantly positive topographical sensitivity ($p < 0.05$, Bonferroni corrected for multiple comparisons). Letters correspond to subgraph identifiers in Fig. S2. Gray area corresponds to the 95% confidence interval of topographical sensitivity across all surrogate subgraphs. (*C*) Distribution of inter-hemispheric topographical sensitivity for each subgraph – rank ordered from smallest to largest. Compared to the null model, subgraphs 1–2 (blue) exhibited significantly negative topographical sensitivity and subgraphs 18–21 (red) exhibited significantly positive topographical sensitivity ($p < 0.05$, Bonferroni corrected for multiple comparisons). (*B* and *C*) These results suggest that edge co-variances have a robust topographical organization in which the strongest edges may form subgraphs over either short or long geographical distances. (*D* and *E*) Example of intra-hemispheric and inter-hemispheric subgraphs with long-range and short-range topographical sensitivities. The short-range subgraph exemplifies strong edges between nearby brain regions, whereas the long-range subgraph exemplifies strong edges between distant brain regions. Strongest 2.5% of edges are plotted here.

edge weights α across the dynamic functional networks of 20 healthy subjects (Fig. 2).

First, we measure the relationship between the ten-fold cross-validation error and each parameter m, β, α, marginalizing over the remaining parameters (Fig. 2A-C). We observe a significant decrease in cross-validation error as the number of subgraphs increases (Pearson $r = -0.42, p < 1 \times 10^{-16}$), suggesting that subgraphs collectively explain more of the statistical space of dynamic functional networks as the number of decomposed subgraphs increases. We also observe a significant positive relationship between cross-validation error and the temporal sparsity parameter (Pearson $r = 0.36$, $p < 1 \times 10^{-16}$), suggesting that subgraphs become less generalized and more specific to individual time windows as the temporal sparsity parameter increases. Despite finding a significant relationship between cross-validation error and the subgraph regularization parameter (Pearson $r = 0.16, p < 1 \times 10^{-16}$), we observed a weaker correlation than between cross-validation error and temporal sparsity, and between cross-validation error and number of subgraphs. These results suggest that cross-validation error is mostly influenced by the number of subgraphs and by the temporal sparsity parameter.

Importantly, we note a potential degeneracy associated with increasing the dimensionality of the subgraph sub-space m towards the dimensionality of the original network. For large m, we expect that subgraphs would sacrifice functional architecture that generalizes over time for architecture that is specific to single edges or single time windows. To characterize these degeneracies, we measure the percentages of sparse temporal expression coefficients and sparse subgraph edge weights as the number of subgraphs increases (Fig. 2D,G). We find a significant positive relationship between the percent of sparse temporal coefficients and the number of subgraphs (Pearson $r = 0.79$, $p < 1 \times 10^{-16}$), suggesting that for a large number of subgraphs, the temporal expression of subgraphs generally falls to zero and is compensated by brief increases in expression only at specific time points. Similarly, we find a significant positive relationship between the percent of sparse subgraph edge weights and the number of subgraphs (Pearson $r = 0.73, p < 1 \times 10^{-16}$), suggesting that for a large number of subgraphs, the subgraph topology approaches a degeneracy in which edge weights fall to zero with only a few non-zero edges constituting a subgraph.

Next, we characterize the effect of β and α on the percent of sparse temporal coefficients (Fig. 2E,F) and the percent of sparse subgraph edge weights (Fig. 2H,I). As expected, we find a strong relationship between the percent of sparse temporal coefficients and β (Pearson $r = 0.52$, $p < 1 \times 10^{-16}$), and no relationship between the percent of sparse subgraph edge weights and β (Pearson $r = -0.01$, $p = 0.04$). These results suggest that by tuning β we can optimize the generalizability of the subgraph expression without altering the sparsity of the subgraph topology. Similarly, we find a weak relationship between the percent of sparse temporal coefficients and α (Pearson $r = 0.15, p < 1 \times 10^{-16}$), and a strong relationship between the percent of sparse subgraph edge weights and α (Pearson $r = -0.53$, $p < 1 \times 10^{-16}$). These results suggest

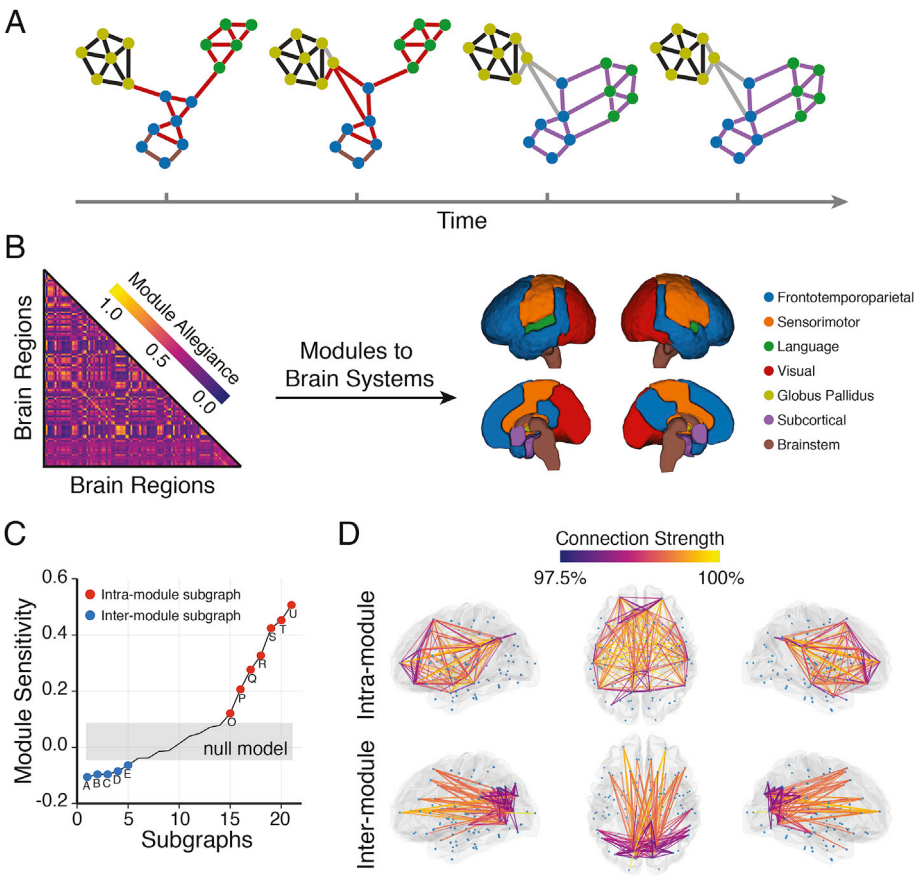

Fig. 4. Network modules and network subgraphs. (*A*) A toy dynamic network that illustrates differential changes to the architecture of network modules and network subgraphs over time. Modules describe a hard-partition of strongly connected network nodes into different groups or "communities" (colored circles) whose composition may change over time. Subgraphs describe a soft-partition of temporally co-varying network edges into different groups (colored lines). Network edges of a single subgraph may occupy a single module (dark blue) or span multiple modules (light blue). (*B*) The module allegiance matrix quantifies the probability that a pair of brain regions occupy the same functional module across all subjects, time windows, and modularity optimizations (*left*; see Methods). Modularity optimization applied to the module allegiance matrix identifies modules that represent putative brain systems (*right*). To quantify the degree to which a subgraph captures the modular organization of the network, we mapped the module allegiance matrix onto the subgraph adjacency matrix using a Pearson correlation similarity metric: positive values indicate greater within-module edge strength, while negative values indicate greater between-module edge strength. (*C*) Distribution of the module sensitivity index of subgraphs in increasing order. Subgraphs exhibited significantly greater module sensitivity in true subgraphs compared to the surrogate data in which NMF was applied to the dynamic network with edges randomly permuted and geometric-relationships between brain regions preserved ($p < 0.05$, Bonferroni corrected for multiple comparisons). Letters correspond to subgraph identifiers in Fig. S2. Gray area corresponds to the 95% confidence interval of module sensitivity across all surrogate subgraphs. Subgraphs with significantly high module sensitivity tend to capture strongest edges within modules (red), and subgraphs with significantly low module sensitivity tend to distribute strongest edges between modules (blue). (*D*) Example of intra-module subgraph *U* and inter-module subgraph *A*. The intra-module subgraph exemplifies strong edges adjoining brain regions of the same module, whereas the inter-module subgraph exemplifies strong edges adjoining brain regions of different modules. Strongest 2.5% of edges are plotted here.

that by tuning α, we can optimize the specificity of the subgraph edge weights without altering the sparsity of the subgraph expression.

Collectively, these results suggest a potential strategy for choosing parameters that achieves a balance between spatial and temporal generalizability and specificity of the subgraphs. Therefore, we average the randomly sampled parameters associated with the lowest 25% cross-validation error. These choices resulted in a number of subgraphs $m = 21$, a temporal sparsity of $\beta = 0.39$, and a subgraph regularization of $\alpha = 0.50$. A plot of all 21 subgraphs can be found in Fig. S2 of the Supplemental Information.

Subgraphs stratify proximal and distributed functional interactions

An open question in network neuroscience is, "What is the spatial organization of naturally formed subgraphs of dynamic brain networks?" We expect that subgraphs stratify coherent groups of functional interactions that reflect more local processing over shorter distances and more distributed processing over longer distances. To investigate the verity of this expectation, we compute the Euclidean distance between all pairs of 112 brain regions defined by the Harvard-Oxford anatomical atlas (Fig. 3A). Based on these distances, we measure the topographical sensitivity for each subgraph based on the correlation between a subgraph's edge weights and the Euclidean distance between the brain regions represented by those edges. To compare topographical sensitivity across subgraphs, we rank subgraphs in increasing order of their correlation with spatial distance, separately for intra-hemisphere and inter-hemisphere edges, and compare the correlations to a null model constructed from geometrically-unconstrained surrogate data. The null model represents the null distribution of correlations when subgraphs are decomposed from dynamic networks in which distance-wise edge relationships have been regressed out.

For intra-hemispheric edges (Fig. 3B), we find that five of twenty-one subgraphs exhibit a more negative correlation than expected by the surrogate model ($p < 0.05$; Bonferroni corrected; subgraphs: P, S, R, Q, H), demonstrating that stronger functional interactions are expressed over shorter distances. Three other subgraphs exhibit a more positive correlation than expected by the surrogate model ($p < 0.05$; Bonferroni corrected; subgraphs: G, A, J), demonstrating that stronger functional interactions are expressed over longer distances. For inter-hemispheric edges (Fig. 3C), we find that two of twenty-one subgraphs demonstrate stronger functional interactions over shorter distances than expected by the surrogate model ($p < 0.05$; Bonferroni corrected, subgraphs: P, C). Four other subgraphs exhibit stronger functional interactions over longer distances than expected by the surrogate model ($p < 0.05$; Bonferroni

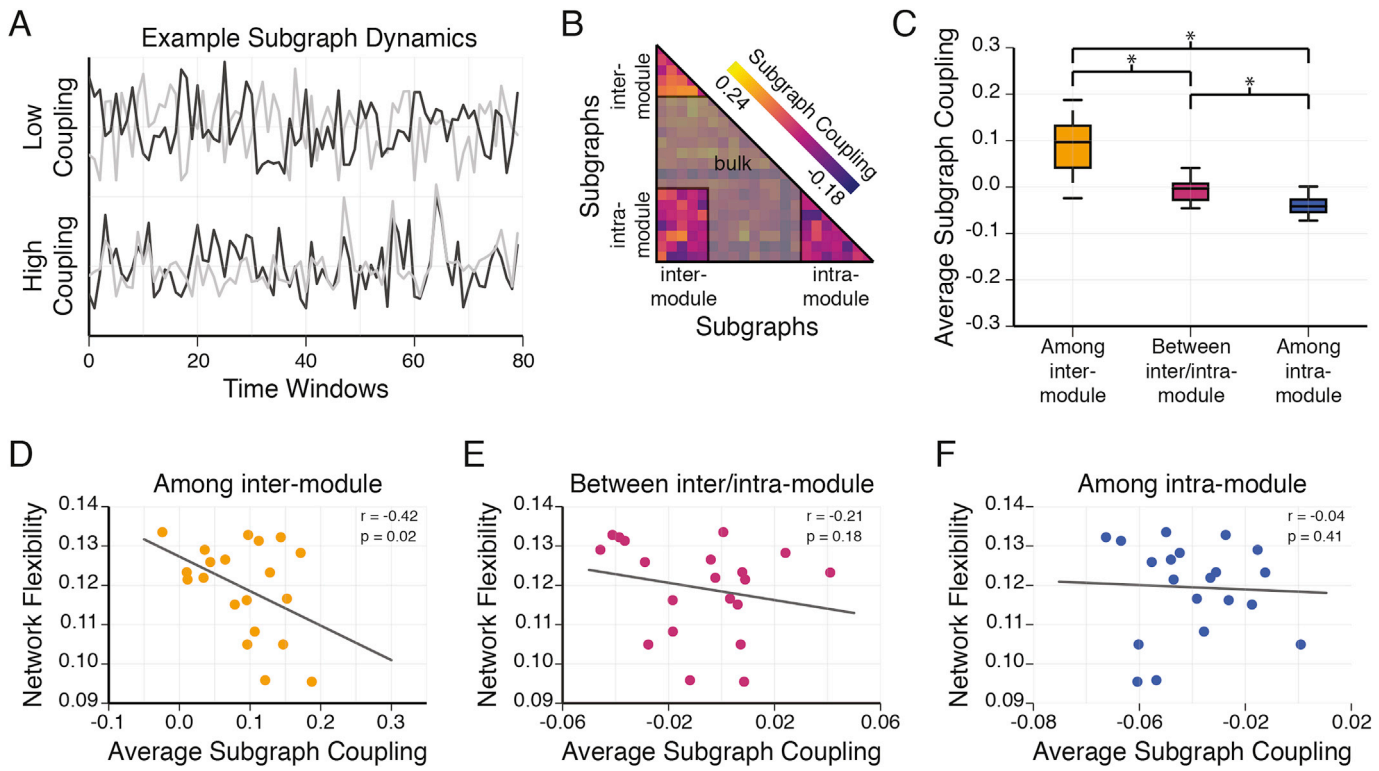

Fig. 5. Dynamical coupling of module sensitive subgraphs. (*A*) We measure the dynamical coupling between subgraphs by computing the Pearson correlation of their time-varying expression coefficients. Shown here are two pairs of example subgraphs that exhibit high coupling (greater correlation) and low coupling (lower correlation). (*B*) Dynamic coupling matrix representing the coupling between every pair of subgraphs, averaged over all subjects and imaging sessions. Rows and columns representing each subgraph are ranked in increasing order of module sensitivity. (*C*) Distribution of subgraph coupling across subjects for three different groups of subgraph-to-subgraph interactions. Coupling among inter-module subgraphs is significantly greater than coupling between inter- and intra-module subgraphs ($t_{19} = 7.45$, $p = 4.8 \times 10^{-7}$) and is significantly greater than coupling among intra-module subgraphs ($t_{19} = 8.9$, $p = 3.1 \times 10^{-8}$). Coupling among intra-module subgraphs is significantly greater than coupling between inter- and intra-module subgraphs ($t_{19} = -4.8$, $p = 1.2 \times 10^{-4}$). The * indicates $p < 0.05$, Bonferroni corrected for multiple comparisons. (*D-F*) Relationship between subgraph coupling of different types and overall network flexibility for each subject – averaged over imaging sessions. We find a significant, inverse relationship between coupling among inter-module subgraphs and network flexibility, suggesting that individuals with tightly-coupled, inter-module subgraphs exhibit a decreased ability to reorganize the modular architecture of their brain network.

corrected; subgraphs: L, B, N, D). The distinction between short-range and long-range subgraphs was similarly observed in subgraphs of a multimodal parcellation of cortex and in subgraphs of a validation dataset consisting of different subjects (see Supplemental Information, Fig. S3A-C, and Fig. S5A-B).

These results suggest that select subgraphs are composed of uniquely local or uniquely distributed functional interactions (see for example short- and long-range, intra- and inter-hemispheric subgraphs (Fig. 3D and E). We also find a second source of variability in the topographical properties of a subgraph that is based on whether adjoining brain regions are located in the same hemisphere or in both hemispheres. In particular, the eight subgraphs that exhibit significant short/long-range sensitivity for intra-hemispheric interactions are largely different than the six subgraphs that exhibit significant short/long-range sensitivity for inter-hemispheric interactions. The exception to this general rule was subgraph *P*, which identifies short-range interactions that span both within and between hemispheres. The finding that subgraphs reveal distributed groups of functional interactions based on physical distance and hemispheres imply that dynamic fluctuations in functional interactions covary differently depending on the geographical location. In light of recent literature (Kopell et al., 2000; Jones et al., 2000; Pinto et al., 2003), these diverse spatial scales of subgraphs may directly relate to the temporal scale of dynamics that the subgraphs support: that is, spatially proximal brain regions can fluctuate in their interaction at a different time-scale than spatially distant brain regions.

Module-based constraints on subgraph architecture

Next we ask, "What constraints might modular organization impart on subgraphs of functional brain networks? Are interactions specified by a subgraph restricted to brain regions of the same module, or can they span multiple modules?" By answering these questions, we can begin piecing together a mechanistic role for subgraphs in the architecture and dynamics of functional brain networks. Based on the perspective that modularity provides a substrate for both integrated and segregated modes of information processing (Bassett et al., 2015; Shine et al., 2016), we hypothesize that the strength of functional interactions will co-vary based on whether they are localized between brain regions within the same module or between brain regions in different modules (Fig. 4A). If such temporal co-variance were to exist within and between modules, functional subgraphs would reflect the various groups of functional interactions that fluctuate over the different time-scales of information processing. We expect that a modular brain network – in which strong functional interactions are more likely contained within a module than between modules – will primarily result in subgraphs with the strongest interactions linking brain regions of the same module. Accordingly, we expect to find fewer subgraphs that capture interactions between modules.

To quantify the sensitivity of functional subgraphs to topology within or between modules, we first assign brain regions into functional modules by computing the module allegiance matrix – a graph representing the probability that any pair of brain regions belong to the same functional module (Fig. 4B; see Methods). Based on the clustering of the module allegiance matrix, we objectively identified seven functional modules that represent the following putative brain systems: fronto-temporoparietal, sensorimotor, language, visual, globus pallidus, subcortical, and brainstem. We next calculate the module sensitivity index that compares each subgraph adjacency matrix to the module

sensitivity matrix (see Methods). Module sensitivity values closer to 1 imply that a subgraph expresses stronger functional interactions between brain regions assigned to the same module, and values closer to−1 imply that a subgraph expresses stronger functional interactions between brain regions assigned to different modules. Probing this index enables us to map subgraph topology to the modular architecture of the network. To identify subgraphs with significant sensitivity to modular architecture, we rank subgraphs in increasing order of their module sensitivity, and we analyze the distribution of sensitivity indices in comparison to a null model built from topologically-unconstrained surrogate data (Fig. 4C). The null model represents the null distribution of correlations when subgraphs are decomposed from dynamic networks in which edges are permuted between brain regions but geometric relationships are preserved.

We find that seven subgraphs exhibit significantly greater module sensitivity than expected in the surrogate model ($p < 0.05$, Bonferroni corrected) – labeled *intra-module* subgraphs. Five other subgraphs exhibit significantly lower module sensitivity than expected in the surrogate model ($p < 0.05$, Bonferroni corrected) – labeled *inter-module* subgraphs. Subgraphs with non-significant module sensitivity were labeled as *bulk* subgraphs. Notably, we observe similar results for subgraphs decomposed from a multimodal parcellation of cortex (see Supplemental Information and Fig. S3D,E). We refer the reader to the Supplemental Information for a more in-depth discussion of several topics, including the relationship between functional subgraphs, the frontotemporoparietal module, and executive/default mode systems. Together, these results suggest that modular architecture heterogenously constrains the functional interactions of subgraphs – some subgraphs express topology that resembles function-specific information processing within modules while other subgraphs express topology that resembles integrative processing across modules.

Module-based constraints on subgraph dynamics

Modular architecture has been shown to compartmentalize neural computations, segregate information processes, and autonomize function (Shine et al., 2016; Bassett et al., 2015) such that dynamics between brain regions in each module may be more correlated than dynamics between regions of different modules. While the relationship between modular architecture and network dynamics has been studied at the scale of individual functional interactions, the influence of a network's modular architecture on the dynamics of a collection of functional interactions, or subgraphs, remains elusive. Intuitively, a subgraph's dynamics are measured by temporal variation in its expression weights – two subgraphs that exhibit more coupled temporal expression profiles have low functional autonomy and two subgraphs that exhibit less coupled temporal expression profiles have high functional autonomy (Fig. 5A; see Methods). Based on the functional role of modules in compartmentalizing information processes, we expect that pairs of subgraphs that link nodes within individual modules might exhibit less coupled dynamics and more functional autonomy than pairs of subgraphs that link nodes between different modules.

To test this hypothesis, we compute the coupling between expression dynamics of each pair of subgraphs – averaged across all subjects and imaging sessions (Fig. 5B). Indeed, we observe high and low coupling between different pairs of subgraphs. To test whether the observed coupling distribution is explained by the different types of module sensitive subgraphs, we compute averaged coupling over all sessions of each subject among inter-module subgraphs, among intra-module subgraphs, and between inter- and intra-module subgraphs (Fig. 5C). Using paired *t*-tests and Bonferroni correction for multiple comparisons, we compare the coupling distributions between different sets of module sensitive subgraphs across subjects and find: (i) inter-module subgraphs are significantly more coupled to one another than they are to intra-module subgraphs ($t_{19} = 7.45$, $p = 4.8 \times 10^{-7}$), (ii) inter-module subgraphs are significantly more coupled to one another than are intra-module subgraphs ($t_{19} = 8.9$, $p = 3.1 \times 10^{-8}$), and (iii) intra-module subgraphs are significantly less coupled to one another than they are to inter-module subgraphs ($t_{19} = -4.8$, $p = 1.2 \times 10^{-4}$). Together, these findings confirm our hypothesis that the degree of dynamical coupling between subgraphs varies based on whether subgraphs map interactions within individual modules or between modules. Subgraphs contained within modules exhibit more autonomous dynamics while subgraphs distributed between modules exhibit less autonomous dynamics. Critically, subgraphs may elucidate coordinated processes related to information integration and segregation at a topologically finer scale beyond what is currently possible with module-level characterization of individual brain regions (Zalesky et al., 2014b; Mattar et al., 2015; Bertolero et al., 2015).

By linking subgraph dynamics to the modular architecture of brain networks, we can ask more fundamentally whether subgraph expression is associated with modular reconfiguration of brain networks. We hypothesized that flexible reconfiguration of network modules – the average rate at which brain regions change their module assignment – is associated with a topological change in the functional interactions that link brain regions of different modules. Intuitively, two brain regions are more likely to occupy the same functional module when the functional interactions between them increases in time – accordingly, reconfiguration across the entire network is more likely when there is a diverse change in the strength of functional interactions between modules. Based on our earlier result that interactions between modules are most effectively captured by inter-module subgraphs, we expect that individuals with reduced coupling between inter-module subgraphs are more likely to show greater module flexibility – signifying a disproportionate increase or decrease in expression of different sets of inter-module subgraphs. We find that coupling among inter-module subgraphs significantly explains individuals' flexible module reconfiguration (Fig. 5D; Pearson's $r = 0.42$, $p = 0.02$). Coupling between inter- and intra-module subgraphs (Fig. 5E; Pearson's $r = -0.21$, $p = 0.18$) and coupling among intra-module subgraphs (Fig. 5F; Pearson's $r = -0.04$, $p = 0.41$) poorly explains inter-individual differences in flexibility of module reconfiguration. Overall, these results suggest that individuals with more synchronized expression between inter-module subgraphs tend to exhibit more rigid and less flexible modular organization. These findings suggest that functional subgraphs that span the network architecture between modules may facilitate the flexibility of an individual's functional brain network.

Are brain regions who are more influential participants in highly coupled subgraphs less likely to switch their module assignment? We next address this question by computing the regional coupling for each brain region – a measure of a brain region's influence in a subgraph weighted by its average expression coupling to all other subgraphs – and the regional flexibility for each brain region – the frequency with which a single brain region changes its module assignment. We find that greater regional coupling is significantly associated with lower regional flexibility (Fig. 6A; Pearson's $r = -0.32$, $p = 4.2 \times 10^{-4}$), suggesting that brain regions that are strongly connected in subgraphs with high coupling to other subgraphs are more rigid in their ability to shift participation between different modules. Practically, we observe that more rigid and highly coupled brain regions tend to overlap with sensory and frontotemporoparietal regions and that more flexible and less coupled brain regions tend to overlap with language and subcortical areas, as well as brainstem (Fig. 6B).

Overall, our findings suggest that the modular architecture of brain networks places fundamental constraints on the various modes of functional interactions between brain regions, as represented by subgraphs. Thus, network modules and network subgraphs may play complimentary roles in guiding dynamics of functional brain networks; network modules prescribe the meso-scale organization of functionally cohesive brain regions, network subgraphs pinpoint the modes in which these brain regions interact with finer granularity. Intriguingly, the dynamics of these

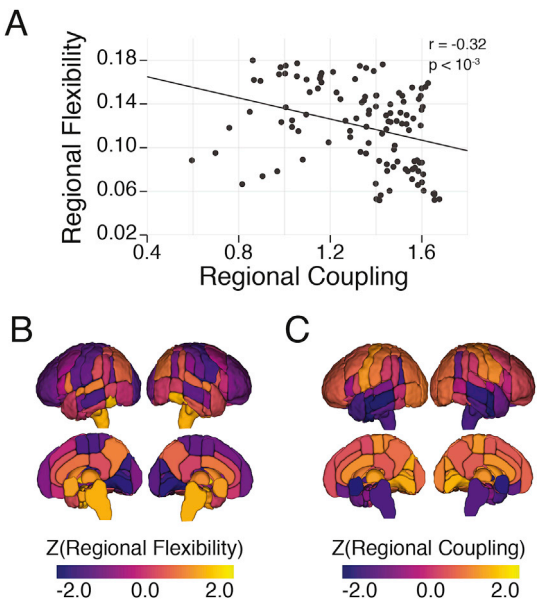

Fig. 6. Subgraph coupling constrains flexibility of individual brain regions. (*A*) Relationship between a brain region's influence on the average coupling of a subgraph to all other subgraphs and the brain region's flexibility to shift its module assignment. We find that brain regions that are more influential participants in tightly-coupled subgraphs are less likely to shift their module assignment and brain regions that are more influential participants in subgraphs that are decoupled from other subgraphs are more likely to shift their module assignment. (*B*) Regional variability of flexibility and coupling across 112 brain regions of the Harvard-Oxford atlas. Frontotemporoparietal and sensory regions exhibit low flexibility and high coupling, while language and subcortical areas, as well as brainstem, tend to be more flexible with low coupling.

individual subgraphs explain fluctuations in integrated and segregated module-level interactions, as well as inter-subject variability in module reorganization.

Discussion

In this study, we put forth a framework for uncovering topological modes of functional interactions from time-evolving brain networks that is based on an unsupervised machine learning tool called non-negative matrix factorization (NMF). We demonstrate that functional brain networks decompose into constituent parts or additive subgraphs that are differentially expressed over time. To effectively identify these subgraphs, we demonstrate an ability to control the number of subgraphs, the temporal sparseness of their expression, and the topological sparseness of their edge weights by manipulating three important parameters of the NMF optimization problem. Using the most robust parameter set for the decomposition, we investigate the topographical and topological basis of the recovered functional subgraphs. We find that subgraphs naturally stratify functional interactions between spatially proximal brain regions and spatially distributed brain regions. Upon closer examination, we observe that these groups of coherent functional interactions underlie a brain architecture that may support integration and segregation of information processing across network modules. Moreover, subgraphs exhibit expression dynamics that reliably signal the modular reorganization of the network.

Machine learning to partition functional brain networks

Network neuroscientists are eager to use computational tools to objectively partition brain networks into naturally organized sub-regions that can deepen our understanding of form and function (Bullmore and Sporns, 2012). Conventional approaches to localize network sub-regions or components are based on graph theoretic tools, such as the minimum-cut procedure (Raj and Wiggins, 2010), that incrementally and continually divide a network along its edges until a chosen criteria is satisfied. While such algorithms can be particularly effective for examining groups of brain regions in anatomically-based structural networks that remain relatively static over short periods of time (Bassett et al., 2010; Hagmann et al., 2008), they are not designed to partition dynamic functional networks whose architecture evolves over time. Dynamic community detection methods fill this void by enabling network scientists to identify modules – cohesive groups of highly interacting brain regions – whose composition could change over time (Bassett et al., 2013a).

Importantly, traditional methods for community detection pursue a *hard* partitioning of brain networks that unambiguously assigns brain regions to a single module based purely on the strength, rather than the topological arrangement, of functional interactions (Porter et al., 2009; Fortunato, 2010). However, recently it has been argued that brain regions may not necessarily have disjoint organization where they fulfill a single functional role. Rather, they may participate in many brain systems along different phases of behavioral and cognitive processing (Wu et al., 2011; Pessoa, 2016; Ding et al., 2016). This fact requires the development of models that are capable of *soft* partitioning the network such that brain regions are allowed to participate in different systems to varying degrees (Yeo et al., 2014; Ball et al., 2011; Ahn et al., 2010). In the brain, such soft partitions can come in the form of so-called link communities (de Reus et al., 2014) or hypergraphs (Bassett et al., 2014; Davison et al., 2015, 2016), although neither of these approaches addresses the existence of dynamics in the partition.

To accommodate greater model flexibility in characterizing network organization and dynamics, we turned to soft partitioning approaches – such as NMF – that afford the ability to statistically learn important rules regarding the behavior of a system through observational data. Unlike hard partitioning approaches – such as community detection – that may be tailored to understand different facets of network organization under particular constraints, soft partitioning deconstructs the statistical space of dynamic networks, provided there is working knowledge of the underlying distribution of that space. Put simply, hard partitioning approaches ask, “Are brain networks organized in a particular way?”, while soft partitioning approaches ask, “What are the rules underlying the observed network organization?” Thus, community detection clusters brain regions into modules based on the strength of their interactions, while NMF groups functional interactions into subgraphs based on their cohesive fluctuations over time.

There are several design considerations for regularizing NMF to identify an “optimal” partitioning of the dynamic network. In our factorization, we impose an L1-sparsity (β) to the temporal expression coefficients that promotes sparser expression of subgraphs in time. Intuitively, this implies that a larger L1-sparsity parameter is associated with a greater likelihood of a subgraph having a zero expression coefficient at a point in time. This regularization strategy allows us to find subgraphs on the basis of the dynamical brain state, which is especially useful when applying NMF to temporally continuous data. The alternative regularization strategy would have been to impose L1-sparsity to the subgraph edge weights, which would promote more spatially-specific subgraphs and a greater likelihood of more subgraphs being expressed at a single point in time. By imposing an L1-sparsity to the subgraph expression coefficients, we increase the dynamical interpretation of the subgraphs while reducing the spatial specificity of the subgraphs. Importantly, our results demonstrate that even when L1-sparsity is applied to the subgraph expression coefficients, the subgraph architectures still tend to follow specific topological and topographical rules.

Spatial and temporal constraints on functional interactions

Fundamentally, subgraphs represent constraints on the functional interactions that the observed network is able to support. Intriguingly, prior work demonstrates that functional subgraphs share resounding similarity to empirically-defined brain systems (Chai et al., 2016a) – such

as the executive system (Shanmugan et al., 2016; Satterthwaite et al., 2013; Nowrangi et al., 2014) and default mode network (Raichle, 2015; Raichle and Snyder, 2007) – supporting the theory that subgraphs prescribe a foundational substrate for the functional interactions underlying cognitive processes and resulting behavior. In our study, we observe that functional subgraphs support a heterogeneous topography of interactions over short and long distances, across and within hemispheres. The observed stratification of distance-based interactions over several subgraphs suggests that these subgraphs also represent unique topological relationships. Our findings are supported by previous studies that use both descriptive statistics (Lohse et al., 2014; Bassett et al., 2010; Klimm et al., 2014) and generative models (Vertes et al., 2012; Betzel et al., 2016) to investigate the relative importance of topography and topology in anatomical networks. Functionally, subgraphs that separately support short- and long-range interactions could be further explored in their potential role in shaping human intelligence, which has been linked to long-distance interactions that promote global efficiency (Li et al., 2009; van den Heuvel et al., 2009; Santarnecchi et al., 2014).

A long purported role of multi-scale network topology is to support segregated processing of functionally specialized information and integrated processing of distinct pieces of information (Bullmore and Sporns, 2012; Betzel and Bassett, 2016). Prior studies demonstrate an ability to capture the functional substrates of integration and segregation during learning (Bassett et al., 2011, 2015; Gerraty et al., 2016), linguistic processing (Chai et al., 2016b; Doron et al., 2012), aging (Chan et al., 2014; Betzel et al., 2015; Meunier et al., 2009), executive function (Braun et al., 2015, 2016), attention (Mattar et al., 2015; Telesford et al., 2016), and rest (Bertolero et al., 2015) using modular decompositions. While modular organization helps compartmentalize function-specific computations within individual modules, the finer-scale topology that actuates this computation has remained elusive. Our results demonstrate that subgraph topologies differentially express functional interactions within and between modules. These within- and between-module subgraphs may help define a clearer functional role for putative provincial and connector hubs in the context of local and distributed computations (Bertolero et al., 2015). In addition to a topological basis for module-based processing, our results suggest that subgraph dynamics may distinguish different stages of integrated and segregated processing based on their relative coupling to one another. The ability to track the relative expression of subgraphs during these stages is crucial for understanding the temporal dynamics that embody so called *cognitive control* processes (Miller and Cohen, 2001) that govern our ability to modulate attention (Kok et al., 2006; Gazzaley and Nobre, 2012; Corbetta and Shulman, 2002) and switch between tasks (Ruge et al., 2013; Kiesel et al., 2010; Vandierendonck et al., 2010).

Network flexibility, the ability for modular architecture to reorganize, has recently been considered a putative functional driver of cognitive control (Cocchi et al., 2013; Bassett et al., 2015; Braun et al., 2015), and is demonstrably altered in psychiatric disorders such as schizophrenia (Braun et al., 2016), which are also characterized by deficits in executive function (Shanmugan et al., 2016; Bassett et al., 2009). In this study, we explore possible relationships between modular reorganization and subgraph dynamics and find that subgraphs that are more sensitive to topology linking different modules together have a strong relationship to network flexibility. First, when individuals express more tightly-coupled dynamics among these subgraphs, they tend to exhibit less flexible functional modules, revealing a more influential role of edge dynamics as a stabilizer of network architecture. Brain regions most likely to embody the role of a network stabilizer are associated with executive and sensory systems – corroborating prior work that finds subgraphs of the executive system are more energetically expressed in young adults compared to children (Chai et al., 2016a), which might be explained by less impulsive and more proactive behavior (Thompson-Schill et al., 2009; Chrysikou et al., 2011). Second, individuals that express highly decoupled between-module subgraphs tend to exhibit more flexible functional modules, instantiating increased autonomy among different sets of module-level processes – a capability that is vital for learning (Bassett et al., 2015).

Methodological considerations

As mentioned earlier in the *Results* section, we provide test-retest reliability results both within our main dataset, as well as in a replication dataset composed of different individuals and acquired on a different scanner at a different research site. However, it is important to note that there is no well accepted approach to confirm reliability of graph sets across iterative measurement. Indeed, a critical challenge is mapping a correspondence between clusters identified in different datasets and algorithm runs. Here, we assess reliability by adopting the Hungarian Assignment Algorithm (Kuhn, 1955), which has been recently used in applications of machine learning on graph sets (Eavani et al., 2015; Leonardi et al., 2014). While we present statistical tests that have been similarly used in these recent studies, future efforts would benefit from further work by statisticians with interests in graph representations of imaging data. In addition, while we find the subgraphs statistically similar across two halves of the main dataset, and between the main dataset and the replication dataset, the observed similarity is not exact and the percent variance explained is not 100. While the unexplained variance could be due to measurement noise, it could also be due to important variations across cohorts. Indeed, it will be interesting in future work to examine whether there are statistically significant differences in the subgraphs detected in different clinical groups, at different ages, or under different task contexts.

Conclusions

In this study, we introduce a machine learning approach for decomposing dynamic functional networks into subgraphs and characterizing their architecture in the context of topography and topology. We show that subgraphs stratify groups of functional interactions that are expressed over a variety of spatial distances. Furthermore, subgraphs express topologies that obey the modular architecture of functional brain networks. These topological constraints lend subgraphs an important ability to signal dynamical states of inter- and intra-module processing. In general, our framework can be used to examine the functional substrates underlying various behavioral and cognitive states, such as the coordination of different groups of functional interactions during a task. Importantly, the NMF-based approach would highlight which brain regions are engaged during these states and how these engagements shift over time. Further methodological development could highlight sequences of subgraph expression that could be used to understand dynamics associated with cognitively effortful tasks, and their alteration in patients with neurological disorders or psychiatric disease, and to pinpoint causal drivers of network dynamics (Mantzaris et al., 2013; Gratton et al., 2014).

Acknowledgments

A.N.K, M.G.M, and D.S.B would like to acknowledge support from the John D. and Catherine T. MacArthur Foundation, the Alfred P. Sloan Foundation, the Army Research Laboratory and the Army Research Office through contract numbers W911NF-10-2-0022 and W911NF-14-1-0679, the National Institutes of Health (2-R01-DC-009209-11, 1R01HD086888-01, R01-MH107235, R01-MH107703, R01MH109520, 1R01NS099348 and R21-M MH-106799), the Office of Naval Research, and the National Science Foundation (BCS-1441502, CAREER PHY-1554488, BCS-1631550, and CNS-1626008).The content is solely the responsibility of the authors and does not necessarily represent the official views of any of the funding agencies.

Appendix A. Supplementary data

Supplementary data related to this article can be found at https://doi.

org/10.1016/j.neuroimage.2017.11.015

References

Ahn, Y.Y., Bagrow, J.P., Lehmann, S., 2010. Link communities reveal multiscale complexity in networks. Nature 466 (7307), 761–764.

Alexander, G.E., Crutcher, M.D., DeLong, M.R., 1990. Basal ganglia-thalamocortical circuits: parallel substrates for motor, oculomotor, "prefrontal" and "limbic" functions. Prog. Brain Res. 85, 119–146.

Alexander-Bloch, A.F., Vértes, P.E., Stidd, R., Lalonde, F., Clasen, L., Rapoport, J., Giedd, J., Bullmore, E.T., Gogtay, N., 2012. The anatomical distance of functional connections predicts brain network topology in health and schizophrenia. Cereb. cortex 23 (1), 127–138.

Ashourvan, A., Gu, S., Mattar, M.G., Vettel, J.M., Bassett, D.S., 2017. The energy landscape underpinning module dynamics in the human brain connectome. NeuroImage 157, 364–380.

Ball, B., Karrer, B., Newman, M.E., 2011. Efficient and principled method for detecting communities in networks. Phys. Rev. E Stat. Nonlin Soft Matter Phys. 84 (3 Pt 2), 036103.

Bassett, D.S., Bullmore, E., 2006. Small-world brain networks. Neuroscientist 12 (6), 512–523.

Bassett, D.S., Bullmore, E.T., Meyer-Lindenberg, A., Apud, J.A., Weinberger, D.R., Coppola, R., 2009. Cognitive fitness of cost-efficient brain functional networks. Proc. Natl. Acad. Sci. U. S. A. 106 (28), 11747–11752.

Bassett, D.S., Greenfield, D.L., Meyer-Lindenberg, A., Weinberger, D.R., Moore, S.W., Bullmore, E.T., 2010. Efficient physical embedding of topologically complex information processing networks in brains and computer circuits. PLoS Comput. Biol. 6 (4), e1000748.

Bassett, D.S., Wymbs, N.F., Porter, M.A., Mucha, P.J., Carlson, J.M., Grafton, S.T., 2011. Dynamic reconfiguration of human brain networks during learning. Proc. Natl. Acad. Sci. U. S. A. 108 (18), 7641–7646.

Bassett, D.S., Porter, M.A., Wymbs, N.F., Grafton, S.T., Carlson, J.M., Mucha, P.J., 2013a. Robust detection of dynamic community structure in networks. Chaos (Woodbury, N.Y.) 23 (1), 013142.

Bassett, D.S., Wymbs, N.F., Rombach, M.P., Porter, M.A., Mucha, P.J., Grafton, S.T., 2013b. Task-based core-periphery organization of human brain dynamics. PLoS Comput. Biol. 9 (9), e1003171.

Bassett, D.S., Wymbs, N.F., Porter, M.A., Mucha, P.J., Grafton, S.T., 2014. Cross-linked structure of network evolution. Chaos 24 (1), 013112.

Bassett, D.S., Yang, M., Wymbs, N.F., Grafton, S.T., 2015. Learning-induced autonomy of sensorimotor systems. Nat. Neurosci. 18 (5), 744–751.

Behzadi, Y., Restom, K., Liau, J., Liu, T.T., 2007. A component based noise correction method (compcor) for bold and perfusion based fmri. Neuroimage 37 (1), 90–101.

Bergstra, J., Bengio, Y., 2012. Random search for hyper-parameter optimization. J. Mach. Learn. Res. 13, 281–305.

Bertolero, M.A., Yeo, B.T.T., D'Esposito, M., 2015. The modular and integrative functional architecture of the human brain. Proc. Natl. Acad. Sci. 112 (49), 6798–6807.

Bertolero, M., Yeo, B., D'Esposito, M., 2017. The Diverse Club: the Integrative Core of Complex Networks arXiv preprint arXiv:1701.01150.

Betzel, R.F., Bassett, D.S., 2016. Multi-scale brain networks. Neuroimage 160, 73–83.

Betzel, R.F., Mišić, B., He, Y., Rumschlag, J., Zuo, X.-N., Sporns, O., 2015. Functional Brain Modules Reconfigure at Multiple Scales across the Human Lifespan, pp. 1–56. Submitted.

Betzel, R.F., Avena-Koenigsberger, A., Goñi, J., He, Y., de Reus, M.A., Griffa, A., Vértes, P.E., Mišic, B., Thiran, J.P., Hagmann, P., van den Heuvel, M., Zuo, X.N., Bullmore, E.T., Sporns, O., 2016. Generative models of the human connectome. NeuroImage 124, 1054–1064.

Braun, U., Schafer, A., Walter, H., Erk, S., Romanczuk-Seiferth, N., Haddad, L., Schweiger, J.I., Grimm, O., Heinz, A., Tost, H., Meyer-Lindenberg, A., Bassett, D.S., 2015. Dynamic reconfiguration of frontal brain networks during executive cognition in humans. Proc. Natl. Acad. Sci. U. S. A. 112 (37), 11678–11683.

Braun, U., Schafer, A., Bassett, D.S., Rausch, F., Schweiger, J.I., Bilek, E., Erk, S., Romanczuk-Seiferth, N., Grimm, O., Geiger, L.S., Haddad, L., Otto, K., Mohnke, S., Heinz, A., Zink, M., Walter, H., Schwarz, E., Meyer-Lindenberg, A., Tost, H., 2016. Dynamic brain network reconfiguration as a potential schizophrenia genetic risk mechanism modulated by NMDA receptor function. Proc. Natl. Acad. Sci. U. S. A. 113 (44), 12568–12573.

Bullmore, E.T., Bassett, D.S., 2011. Brain graphs: graphical models of the human brain connectome. Annu. Rev. Clin. Psychol. 7 (1), 113–140.

Bullmore, E.T., Sporns, O., 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. Nat. Rev. Neurosci. 10 (3), 186–198.

Bullmore, E., Sporns, O., 2012. The economy of brain network organization. Nat. Rev. Neurosci. 13 (5), 336–349.

Calhoun, V.D., Miller, R., Pearlson, G., Adali, T., 2014. The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. Neuron 84 (2), 262–274.

Chai, X.J., Castañón, A.N., Öngür, D., Whitfield-Gabrieli, S., 2012. Anticorrelations in resting state networks without global signal regression. Neuroimage 59 (2), 1420–1428.

Chai, L.R., Khambhati, A.N., Ciric, R., Moore, T., Gur, R.C., Gur, R.E., Satterthwaite, T.D., Bassett, D.S., 2016a. Evolution of brain network dynamics in neurodevelopment. Netw. Neurosci. 1 (1), 14–30.

Chai, L.R., Mattar, M.G., Blank, I.A., Fedorenko, E., Bassett, D.S., 2016b. Functional network dynamics of the language system. Cerebr. Cortex 26 (11), 4148–4159.

Chan, M.Y., Park, D.C., Savalia, N.K., Petersen, S.E., Wig, G.S., 2014. Decreased segregation of brain systems across the healthy adult lifespan. Proc. Natl. Acad. Sci. 111 (46), 4997–5006.

Chrysikou, E.G., Novick, J.M., Trueswell, J.C., Thompson-Schill, S.L., 2011. The other side of cognitive control: can a lack of cognitive control benefit language and cognition? Top. Cogn. Sci. 3 (2), 253–256.

Cocchi, L., Zalesky, A., Fornito, A., Mattingley, J.B., 2013. Dynamic cooperation and competition between brain systems during cognitive control. Trends Cognit. Sci. 17 (10), 493–501.

Corbetta, M., Shulman, G.L., 2002. Control of goal-directed and stimulus-driven attention in the brain. Nat. Rev. Neurosci. 3 (3), 201–215.

Davison, E.N., Schlesinger, K.J., Bassett, D.S., Lynall, M.E., Miller, M.B., Grafton, S.T., Carlson, J.M., 2015. Brain network adaptability across task states. PLoS Comput. Biol. 11 (1), e1004029.

Davison, E.N., Turner, B.O., Schlesinger, K.J., Miller, M.B., Grafton, S.T., Bassett, D.S., Carlson, J.M., 2016. Individual differences in dynamic functional brain connectivity across the human lifespan. PLoS Comput. Biol. 12 (11), e1005178.

de Reus, M.A., van den Heuvel, M.P., 2013. Rich club organization and intermodule communication in the cat connectome. J. Neurosci. 33 (32), 12929–12939.

de Reus, M.A., Saenger, V.M., Kahn, R.S., van den Heuvel, M.P., 2014. An edge-centric perspective on the human connectome: link communities in the brain. Philos. Trans. R. Soc. Lond B Biol. Sci. 369, 1653.

Deco, G., Jirsa, V.K., 2012. Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. J. Neurosci. 32 (10), 3366–3375.

Deco, G., Senden, M., Jirsa, V., 2012. How anatomy shapes dynamics: a semi-analytical study of the brain at rest by a simple spin model. Front. Comput. Neurosci. 6, 68.

Ding, Z., Zhang, X., Sun, D., Luo, B., 2016. Overlapping community detection based on network decomposition. Sci. Rep. 6, 24115.

Doron, K.W., Bassett, D.S., Gazzaniga, M.S., 2012. Dynamic network structure of interhemispheric coordination. Proc. Natl. Acad. Sci. U. S. A. 109 (46), 18661–18668.

Eavani, H., Satterthwaite, T.D., Filipovych, R., Gur, R.E., Gur, R.C., Davatzikos, C., 2015. Identifying sparse connectivity patterns in the brain using resting-state fMRI. Neuroimage 105, 286–299.

Van Essen, D.C., Anderson, C.H., Felleman, D.J., 1992. Information processing in the primate visual system: an integrated systems perspective. Science 255 (5043), 419–423.

Felleman, D.J., Van Essen, D.C., 1991. Distributed hierarchical processing in the primate cerebral cortex. Cereb. Cortex 1 (1), 1–47.

Fornito, A., Zalesky, A., Breakspear, M., 2015. The connectomics of brain disorders. Nat. Rev. Neurosci. 16 (3), 159–172.

Fortunato, S., 2010. Community detection in graphs. Phys. Rep. 486 (3–5), 75–174.

Friston, K.J., Williams, S., Howard, R., Frackowiak, R.S., Turner, R., 1996. Movement-related effects in fmri time-series. Magn. Reson. Med. 35 (3), 346–355.

Gazzaley, A., Nobre, A.C., 2012. Top-down modulation: bridging selective attention and working memory. Trends Cogn. Sci. 16 (2), 129–135.

Gerraty, R.T., Davidow, J.Y., Foerde, K., Galvan, A., Bassett, D.S., Shohamy, D., 2016. Dynamic Flexibility in Striatal-cortical Circuits Supports Reinforcement Learning bioRxiv 094383, 094383.

Glasser, M.F., Coalson, T.S., Robinson, E.C., Hacker, C.D., Harwell, J., Yacoub, E., Ugurbil, K., Andersson, J., Beckmann, C.F., Jenkinson, M., et al., 2016. A multi-modal parcellation of human cerebral cortex. Nature 536 (7615), 171–178.

Good, B.H., de Montjoye, Y.A., Clauset, A., 2010. Performance of modularity maximization in practical contexts. Phys. Rev. E Stat. Nonlin Soft Matter Phys. 81 (4 Pt 2), 046106.

Gratton, C., Lee, T.G., Nomura, E.M., D'Esposito, M., 2014. Perfusion MRI indexes variability in the functional brain effects of theta-burst transcranial magnetic stimulation. PLoS One 9 (7), e101430.

Greene, D., 2009. A matrix factorization approach for integrating multiple data views. In: Joint European Conference on Machine Learning and Knowledge Discovery in Databases. S. B, Heidelberg.

Greene, D., Cagney, G., Krogan, N., Cunningham, P., 2008. Ensemble non-negative matrix factorization methods for clustering protein-protein interactions. Bioinformatics 24 (15), 1722–1728.

Gu, S., Betzel, R.F., Mattar, M.G., Cieslak, M., Delio, P.R., Grafton, S.T., Pasqualetti, F., Bassett, D.S., 2017. Optimal trajectories of brain state transitions. NeuroImage 148, 305–317.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C.J., Wedeen, V.J., Sporns, O., 2008. Mapping the structural core of human cerebral cortex. PLoS Biol. 6 (7), e159.

Hindriks, R., Adhikari, M.H., Murayama, Y., Ganzetti, M., Mantini, D., Logothetis, N.K., Deco, G., 2016. Can sliding-window correlations reveal dynamic functional connectivity in resting-state fMRI? NeuroImage 127, 242–256.

Hutchison, R.M., Womelsdorf, T., Allen, E.A., Bandettini, P.A., Calhoun, V.D., Corbetta, M., Della Penna, S., Duyn, J.H., Glover, G.H., Gonzalez-Castillo, J., Handwerker, D.A., Keilholz, S., Kiviniemi, V., Leopold, D.A., de Pasquale, F., Sporns, O., Walter, M., Chang, C., 2013. Dynamic functional connectivity: promise, issues, and interpretations. NeuroImage 80, 360–378.

Jenkinson, M., 2004. Improving the registration of b0-distorted epi images using calculated cost function weights. In: Tenth International Conference on Functional Mapping of the Human Brain.

Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage 17 (2), 825–841.

Jeub, L.G.S., Bazzi, M., Jutla, I.S., Mucha, P.J., 2016. A Generalized Louvain Method for Community Detection Implemented in MATLAB.

Jo, H.J., Saad, Z.S., Simmons, W.K., Milbury, L.A., Cox, R.W., 2010. Mapping sources of correlation in resting state fmri, with artifact detection and removal. Neuroimage 52 (2), 571–582.
Jones, S.R., Pinto, D.J., Kaper, T.J., Kopell, N., 2000. Alpha-frequency rhythms desynchronize over long cortical distances: a modeling study. J. Comput. Neurosci. 9 (3), 271–291.
Kaiser, M., 2011. A tutorial in connectome analysis: topological and spatial features of brain networks. Neuroimage 57 (3), 892–907.
Khambhati, A.N., Davis, K.A., Oommen, B.S., Chen, S.H., Lucas, T.H., Litt, B., Bassett, D.S., dec 2015. Dynamic network drivers of seizure generation, propagation and termination in human neocortical epilepsy. PLOS Comput. Biol. 11 (12), e1004608.
Khambhati, A.N., Bassett, D.S., Oommen, B.S., Chen, S.H., Lucas, T.H., Davis, K.A., Litt, B., 2017. Recurring functional interactions predict network architecture of interictal and ictal states in neocortical epilepsy. eNeuro 4 (1) pp.ENEURO-0091.
Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A.M., Koch, I., 2010. Control and interference in task switching–a review. Psychol. Bull. 136 (5), 849–874.
Kim, J., Park, H., 2011. Fast nonnegative matrix factorization: an active-set-like method and comparisons. SIAM J. Sci. Comput. 33 (6), 3261–3281.
Kim, J., He, Y., Park, H., 2014. Algorithms for Nonnegative Matrix and Tensor Factorizations: a Unified View Based on Block Coordinate Descent Framework, vol 58.
Klimm, F., Bassett, D.S., Carlson, J.M., Mucha, P.J., 2014. Resolving structural variability in network models and the brain. PLoS Comput. Biol. 10 (3), e1003491.
Kok, A., Ridderinkhof, K.R., Ullsperger, M., 2006. The control of attention and actions: current research and future developments. Brain Res. 1105 (1), 1–6.
Kopell, N., Ermentrout, G.B., Whittington, M.A., Traub, R.D., 2000. Gamma rhythms and beta rhythms have different synchronization properties. Proc. Natl. Acad. Sci. U. S. A. 97 (4), 1867–1872.
Kopell, N.J., Gritton, H.J., Whittington, M.A., Kramer, M.A., 2014. Beyond the connectome: the dynome. Neuron 83 (6), 1319–1328.
Kuhn, H.W., 1955. The Hungarian method for the assignment problem. Nav. Res. Logist. (NRL) 2 (1–2), 83–97.
Lee, D.D., Seung, H.S., Seung, S., 1999. Learning the parts of objects by non-negative matrix factorization. Nature 401 (6755), 788–791.
Leonardi, N., Richiardi, J., Gschwind, M., Simioni, S., Annoni, J.-M., Schluep, M., Vuilleumier, P., Ville, D.V.D., dec 2013. Principal components of functional connectivity : a new approach to study dynamic brain connectivity during rest. NeuroImage 83, 937–950.
Leonardi, N., Shirer, W.R., Greicius, M.D., Van De Ville, D., jul 2014. Disentangling dynamic networks: separated and joint expressions of functional connectivity patterns in time. Hum. Brain Mapp. 5995, 5984–5995.
Li, Y., Liu, Y., Li, J., Qin, W., Li, K., Yu, C., Jiang, T., 2009. Brain anatomical network and intelligence. PLoS Comput. Biol. 5 (5), e1000395.
Lohse, C., Bassett, D.S., Lim, K.O., Carlson, J.M., 2014. Resolving anatomical and functional structure in human brain organization: identifying mesoscale organization in weighted network representations. PLoS Comput. Biol. 10 (10), e1003712.
Mantzaris, A.V., Bassett, D.S., Wymbs, N.F., Estrada, E., Porter, M.A., Mucha, P.J., Grafton, S.T., Higham, D.J., 2013. Dynamic network centrality summarizes learning in the human brain. J. Complex Netw. 1 (1), 83–92.
Mattar, M.G., Cole, M.W., Thompson-Schill, S.L., Bassett, D.S., 2015. A functional cartography of cognitive systems. PLoS Comput. Biol. 11 (12), e1004533.
Mattar, M.G., Thompson-Schill, S.L., Bassett, D.S., 2016a. The network architecture of value learning. Netw. Neurosci. 1–27, 0, ja.
Mattar, M.G., Wymbs, N.G., Bock, A.S., Aguirre, G.K., Grafton, S.T., Bassett, D.S., 2016b. Predicting future learning from baseline network architecture. bioRxiv 10 (1101), 056861.
Medaglia, J.D., Lynall, M.-E., Bassett, D.S., 2015. Cognitive network neuroscience. J. Cognit. Neurosci. 27 (8), 1471–1491.
Meunier, D., Achard, S., Morcom, A., Bullmore, E., 2009. Age-related changes in modular organization of human brain functional networks. Neuroimage 44 (3), 715–723.
Miller, E.K., Cohen, J.D., 2001. An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. 24, 167–202.
Monti, S., Tamayo, P., Mesirov, J., Golub, T., Sebastiani, P., Kohane, I.S., Ramoni, M.F., 2003. Consensus clustering: a resampling-based method for class discovery and visualization of gene expression microarray data. Mach. Learn. 52 (1–2), 91–118.
Mucha, P.J., Richardson, T., Macon, K., Porter, M.A., Onnela, J.-P.J.-P., 2010. Community structure in time-dependent, multiscale, and multiplex networks. Sci. (New York, N. Y.) 328 (5980), 876–878.
Murphy, K., Birn, R.M., Handwerker, D.A., Jones, T.B., Bandettini, P.A., 2009. The impact of global signal regression on resting state correlations: are anti-correlated networks introduced? Neuroimage 44 (3), 893–905.
Newman, M.E., 2006. Modularity and community structure in networks. Proc. Natl. Acad. Sci. 103 (23), 8577–8582.
Newman, M.E.J., Girvan, M., 2004. Finding and evaluating community structure in networks. Phys. Rev. E - Stat. Nonlinear, Soft Matter Phys. 69 (2 2), 1–15.
Nowrangi, M.A., Lyketsos, C., Rao, V., Munro, C.A., 2014. Systematic review of neuroimaging correlates of executive functioning: converging evidence from different clinical populations. J. Neuropsychiatry Clin. Neurosci. 26 (2), 114–125.
Park, H.J., Friston, K., 2013. Structural and functional brain networks: from connections to cognition. Sci. (New York, N. Y.) 342 (6158), 1238411.
Pessoa, L., 2016. Beyond disjoint brain networks: overlapping networks for cognition and emotion. Behav. Brain Sci. 9.
Pinto, D.J., Jones, S.R., Kaper, T.J., Kopell, N., 2003. Analysis of state-dependent transitions in frequency and long-distance coordination in a model oscillatory cortical circuit. J. Comput. Neurosci. 15 (2), 283–298.
Porter, M. a., Onnela, J.-P., Mucha, P.J., 2009. Communities in networks. Am. Math. Soc. 56 (9), 0–26.
Prieto, G. a., Parker, R.L., Vernon, F.L., 2009. A Fortran 90 library for multitaper spectrum analysis. Comput. Geosci. 35 (8), 1701–1710.
Raichle, M.E., 2015. The brain's default mode network. Annu. Rev. Neurosci. 38, 433–447.
Raichle, M.E., Snyder, A.Z., 2007. A default mode of brain function: a brief history of an evolving idea. Neuroimage 37 (4), 1083–1090.
Raj, A., Wiggins, C.H., 2010. An information-theoretic derivation of min-cut-based clustering. IEEE Trans. Pattern Anal. Mach. Intell. 32 (6), 988–995.
Reijneveld, J.C., Ponten, S.C., Berendse, H.W., Stam, C.J., 2007. The application of graph theoretical analysis to complex networks in the brain. Clin. Neurophysiol. 118 (11), 2317–2331.
Roberts, J.A., Perry, A., Lord, A.R., Roberts, G., Mitchell, P.B., Smith, R.E., Calamante, F., Breakspear, M., 2016. The contribution of geometry to the human connectome. NeuroImage 124, 379–393.
Rubinov, M., Bassett, D.S., 2011. Emerging evidence of connectomic abnormalities in schizophrenia. J. Neurosci. 31 (17), 6263–6265.
Rubinov, M., Sporns, O., 2010. Complex network measures of brain connectivity: uses and interpretations. Neuroimage 52 (3), 1059–1069.
Ruge, H., Jamadar, S., Zimmermann, U., Karayanidis, F., 2013. The many faces of preparatory control in task switching: reviewing a decade of fMRI research. Hum. Brain Mapp. 34 (1), 12–35.
Saad, Z.S., Gotts, S.J., Murphy, K., Chen, G., Jo, H.J., Martin, A., Cox, R.W., 2012. Trouble at rest: how correlation patterns and group differences become distorted after global signal regression. Brain connect. 2 (1), 25–32.
Santarnecchi, E., Galli, G., Polizzotto, N.R., Rossi, A., Rossi, S., Feb 2014. Efficiency of weak brain connections support general cognitive functioning. Hum. Brain Mapp. 35 (9), 4566–4582.
Satterthwaite, T.D., Wolf, D.H., Erus, G., Ruparel, K., Elliott, M.A., Gennatas, E.D., Hopson, R., Jackson, C., Prabhakaran, K., Bilker, W.B., Calkins, M.E., Loughead, J., Smith, A., Roalf, D.R., Hakonarson, H., Verma, R., Davatzikos, C., Gur, R.C., Gur, R.E., 2013. Functional maturation of the executive system during adolescence. J. Neurosci. 33 (41), 16249–16261.
Schlesinger, K.J., Turner, B.O., Lopez, B.A., Miller, M.B., Carlson, J.M., 2017 Feb 1. Age-dependent changes in task-based modular organization of the human brain. Neuroimage 146, 741–762.
Senden, M., Goebel, R., Deco, G., 2012. Structural connectivity allows for multi-threading during rest: the structure of the cortex leads to efficient alternation between resting state exploratory behavior and default mode processing. Neuroimage 60 (4), 2274–2284.
Senden, M., Deco, G., de Reus, M.A., Goebel, R., van den Heuvel, M.P., 2014. Rich club organization supports a diverse set of functional network configurations. Neuroimage 96, 174–182.
Shanmugan, S., Wolf, D.H., Calkins, M.E., Moore, T.M., Ruparel, K., Hopson, R.D., Vandekar, S.N., Roalf, D.R., Elliott, M.A., Jackson, C., Gennatas, E.D., Leibenluft, E., Pine, D.S., Shinohara, R.T., Hakonarson, H., Gur, R.C., Gur, R.E., Satterthwaite, T.D., 2016. Common and dissociable mechanisms of executive system dysfunction across psychiatric disorders in youth. Am. J. Psychiatry 173 (5), 517–526.
Shine, J.M., Bissett, P.G., Bell, P.T., Koyejo, O., Balsters, J.H., Gorgolewski, K.J., Moodie, C.A., Poldrack, R.A., 2016 Oct 19. The dynamics of functional brain networks: integrated network states during cognitive task performance. Neuron 92 (2), 544–554.
Smith, S.M., 2002. Fast robust automated brain extraction. Hum. Brain Mapp. 17 (3), 143–155.
Sporns, O., 2013. The human connectome: origins and challenges. Neuroimage 80, 53–61.
Sporns, O., 2014. Contributions and challenges for network models in cognitive neuroscience. Nat. Neurosci. 17 (5), 652–660.
Sporns, O., Betzel, R.F., 2016. Modular brain networks. Annu. Rev. Psychol. 67, 613–640.
Sporns, O., Tononi, G., Kotter, R., 2005. The human connectome: a structural description of the human brain. PLoS Comput. Biol. 1 (4), e42.
Telesford, Q.K., Simpson, S.L., Burdette, J.H., Hayasaka, S., Laurienti, P.J., 2011. The brain as a complex system: using network science as a tool for understanding the brain. Brain Connect. 1 (4), 295–308.
Telesford, Q.K., Lynall, M.E., Vettel, J., Miller, M.B., Grafton, S.T., Bassett, D.S., 2016. Detection of functional brain network reconfiguration during task-driven cognitive states. Neuroimage 142, 198–210.
Thompson-Schill, S.L., Ramscar, M., Chrysikou, E.G., 2009. Cognition without control: when a little frontal lobe goes a long way. Curr. Dir. Psychol. Sci. 18 (5), 259–263.
van den Heuvel, M.P., Sporns, O., 2011. Rich-club organization of the human connectome. J. Neurosci. 31 (44), 15775–15786.
van den Heuvel, M.P., Stam, C.J., Kahn, R.S., Hulshoff Pol, H.E., 2009. Efficiency of functional brain networks and intellectual performance. J. Neurosci. Offic. J. Soc. Neurosci. 29 (23), 7619–7624.
Vandierendonck, A., Liefooghe, B., Verbruggen, F., 2010. Task switching: interplay of reconfiguration and interference control. Psychol. Bull. 136 (4), 601–626.
Vertes, P.E., Alexander-Bloch, A.F., Gogtay, N., Giedd, J.N., Rapoport, J.L., Bullmore, E.T., 2012. Simple models of human brain functional networks. Proc. Natl. Acad. Sci. U. S. A. 109 (15), 5868–5873.
Wei, M., Qin, J., Yan, R., Bi, K., Liu, C., Yao, Z., Lu, Q., 2017 Apr. Abnormal dynamic community structure of the salience network in depression. J. Magn. Reson. Imag. 45 (4), 1135–1143.
Wu, K., Taki, Y., Sato, K., Sassa, Y., Inoue, K., Goto, R., Okada, K., He, Y., Evans, A.C., Fukuda, H., 2011. The overlapping community structure of structural brain network in young healthy individuals. PLoS One 6 (5).

Wymbs, N.F., Grafton, S.T., 2014. The human motor system supports sequence-specific representations over multiple training-dependent timescales. Cereb. Cortex 25 (11), 4213–4225.
Yeo, B.T., Krienen, F.M., Chee, M.W., Buckner, R.L., 2014. Estimates of segregation and overlap of functional connectivity networks in the human cerebral cortex. Neuroimage 88, 212–227.
Zalesky, A., Fornito, A., Cocchi, L., Gollo, L.L., Breakspear, M., 2014a. Time-resolved resting-state brain networks. Proc. Natl. Acad. Sci. U. S. A. 111 (28), 10341–10346.
Zalesky, A., Fornito, A., Cocchi, L., Gollo, L.L., Breakspear, M., 2014b. Time-resolved resting-state brain networks. Proc. Natl. Acad. Sci. U. S. A. 111 (28), 10341–10346.